arXiv:1603.03999v2 [quant-ph] 22 Nov 2016

The Classification of Stabilizer Operations over Qubits

Daniel Grier* Luke Schaeffer†

Abstract

We present a complete classification of quantum stabilizer gates in terms of the functions they generate assuming the ability to swap qubits and use ancillary workspace. Because we view these stabilizer circuits as subroutines of some general quantum computation, we insist that any ancilla qubits used during the computation must not change in an input-dependent manner. This is the first attempt at a quantum extension of the classification of reversible classical gates introduced by Aaronson et al., another part of an ambitious program to classify all quantum gate sets.

The classification uses, at its center, a reinterpretation of the tableau representation of stabilizer gates to give circuit decompositions, from which elementary generators can easily be extracted. There are a total of 57 different stabilizer classes generated in this way, 30 of which arise from the single-qubit subgroups of the Clifford group. At a high level, the remaining classes are arranged according to the bases they preserve. For instance, the CNOT gate preserves the X and Z bases because it maps X-basis elements to X-basis elements and Z-basis elements to Z-basis elements. The remaining classes are characterized by more subtle tableau invariants; for instance, the T_4 and phase gate generate a proper subclass of Z-preserving gates.

1 Introduction

A common thread through much of quantum computing and indeed through computation in general is the manner in which a few elementary gates often suffice for universal computation. Indeed, this "pervasiveness of universality" as described in recent work of Aaronson, Grier, and Schaeffer [3] helps to explain why a random gate on n qubits suffices for universal quantum computation. In that work, the authors give a complete classification of *classical* reversible gates in terms of the functions over bits they generate. Indeed, they find that even when they restrict themselves to classical gates, a rich structure emerges.

Of course, the ultimate goal would be a complete classification of *quantum* gate sets based on the functions over qubits they generate. Unfortunately, the extraordinary difficulty of this problem given the current state of mathematics renders such a task nearly impossible. Not even a full classification of the subgroups of a three-qubit system is known![1] This might be surprising given how well we understand random instances and even gate sets that contain particular gates such as CNOT [6, 19]. However, a full classification begets a complete understanding of *all* possible behaviors, despite their strangeness or rarity (see, for example, the sporadic gate sets in the lattice

*MIT. Email: grierd@mit.edu. Supported by an NSF Graduate Research Fellowship under Grant No. 1122374.
†MIT. Email: lrs@mit.edu.

[1]The difficulty in classifying the subgroups of $\mathrm{SU}(N)$ arises not from the infinite classes but from the finite ones. In fact, even the finite subgroups of SU(5) remain unclassified. This motivates our focus on finite, discrete classes such as the stabilizer group.

of Aaronson et al. [3]). Nevertheless, there has been some recent and encouraging progress on *some* classification problems, in particular, on classifying Hamiltonians (which can be applied for any period of time), rather than discrete gate sets. For instance, Bouland, Mančinska, and Zhang [7] recently classified all 2-qubit commuting Hamiltonians while Childs et al. [9] characterized all 2-qubit Hamiltonians when restricted to circuits over two qubits. Additionally, Aaronson and Bouland [1] completed a classification for linear optics of 2-mode beamsplitters, which relied heavily on the characterization [15] of the finite subgroups of SU(3), underscoring the difficulty of quantum gate classification.

This paper contributes its own classification of quantum gate sets by giving a complete classification of the so-called stabilizer gates, where we allow the swapping of qubits and the use of ancillary workspace. To provide some context, stabilizer gates[2] are a discrete set of gates generated by the CNOT gate, the Hadamard gate, and the $\frac{\pi}{4}$-phase gate. Stabilizer circuits are somewhat remarkable in that they may in fact be integral to our eventual development of a *general-purpose* quantum computer. Since quantum error correction will likely play a large role in determining when a quantum computer will be viable, there has been considerable research in building and analyzing quantum error correcting codes. The stabilizer formalism arose as powerful way of unifying the analyses of many of these codes [12], and as a consequence, understanding the nature of the stabilizer states has been of particular interest [11, 17].

Our model is motivated by these interests. That is, we view the application of a stabilizer circuit as the subroutine of a general quantum computation. In fact, the stabilizer operations are exactly those operations which can be easily computed transversally in many fault-tolerant schemes of quantum computing (e.g., the Shor code [20] or the [[7,1,3]] Steane code [21]). For this reason, we insist that any ancillary qubits used during the computation return to their initial state by the end of the computation so that the delicate quantum coherence of the protocol is not broken. The ancilla states of our model are very similar to the "reusable magic states" of Anderson [5]. In fact, our classification closes many of the open problems left by that paper. Nevertheless, we regard the classification of stabilizer gates as an interesting mathematical problem in its own right, as it exposes many of the beautiful symmetries inherent in the subgroups of the stabilizer operations.

Because the stabilizer operations on n qubits only generate a discrete, finite set, it has always been widely assumed that they do not suffice for universal quantum computation. Nevertheless, stabilizer gates drew attention when Gottesman and Knill [13] showed that they could be efficiently simulated with a classical computer. This result was surprising to many because entanglement was traditionally thought of as a barrier to classical simulation.

This result was subsequently improved by Aaronson and Gottesman [2], who showed that stabilizer gates have surprisingly little computational power. First notice that the stabilizer gates can be seen as a generalization of the affine reversible gates [3]; CNOT and NOT are stabilizer gates, and hence all affine reversible gates are stabilizer gates. What Aaronson and Gottesman realized is that this is not just a good intuition for the stabilizer gates, but that stabilizer circuits are complete for the class $\oplus\mathsf{L}$ (or *parity*-L), for which the complete problem is the solution of a linear system over $\mathbb{F}_2$. Indeed, $\oplus\mathsf{L}$ is widely regarded to be weaker than all of polynomial time computation. In that work, Aaronson and Gottesman proposed the *tableau representation* of a stabilizer gate, which we reinterpret to serve as one of the principal components of our classification theorem. See Section 4 for those details.

[2]These gates are often called Clifford operations in the literature due to their alternative characterization as normalizers of the Pauli group.

1.1 Results

We present a full inclusion lattice of the 57 different classes of stabilizer gates. Refer to Figure 4 for a diagram of the 30 single-qubit classes and Figure 5 for the remaining lattice of multi-qubit gates. Notation for the generators of the classes depicted in those diagrams is given in Section 3. We list some highlights of the classification below:

(1) **Conserved Quantities.** The invariants that dictate the structure of the lattice are most easily described by appealing to the tableau representation of a stabilizer circuit (see Section 4). Nevertheless, there are some common themes that persist throughout classification. One major theme is that of basis preservation. We say that a gate "preserves" a basis if it acts as an automorphism on the elements of that basis. For instance, the CNOT gate preserves the X and Z bases, but not the Y basis. There are classes which correspond to all 8 possible ways to preserve the X, Y, and Z bases. There are also classes in which each entangling operation only changes the sign of a certain basis, such as the CSIGN gate, or classes in which each entangling operation acts like an orthogonal transformation over the computational basis states, such as the T_4 gate (see Section 3.2). Finally, a class can be constrained by its ability to rotate a single qubit. See Section 6 for a more formal exposition.

(2) **Finite Generation.** Every class can be generated by a single gate on at most four qubits. In fact, given a set of gates generating some class, there always exist three gates from that set that generate the same class. Moreover, the classification implies that the canonical set of stabilizer generators—CNOT, Hadamard, and phase—is *not* a minimal set of generators in our model. Indeed, it turns out that with the aid of ancilla qubits, CNOT and Hadamard generate a phase gate. This is well-known [4, 16], but comes as a simple consequence of our classification theorem.

(3) **Ancillas.** More generally, giving a stabilizer gate access to ancillary qubits often increases the sorts of functions it can compute. A priori, one might suspect that extracting all functionality from a large entangling stabilizer gate would require large highly-entangled ancilla states. Nevertheless, our classification shows that only a constant number of one and two-qubit ancillary states are ever needed. In fact, an even stronger result is true. Namely, our classification holds even when we allow the ancillas to change in an input-independent manner,[3] as would be natural for a stabilizer subroutine applied unconditionally in a quantum circuit. See Section 5 for further discussion.

(4) **Algorithms.** Our classification implies a linear time algorithm which, given the tableau of a gate G, decides which class G belongs to. As a consequence, we can identify the class for a list of tableaux (by running the algorithm for each one, then looking up the least upper bound on our lattice diagram), or determine whether one gate generates another (by running the algorithm for both, and checking whether the one class contains the other). In fact, to witness that G generates some class $\mathcal{C}$ (or higher) in the classification, one only needs to view a constant number of bits of the tableau. These details are discussed in sections 9 and 10.

(5) **Enumeration.** For each class $\mathcal{C}$ and for all n, we give explicit formulas for the number of gates in $\mathcal{C}$ on n-qubits. The enumeration usually leads to efficient canonical forms for circuits

[3] Aaronson et al. called this the "Loose Ancilla Rule," and it *does* affect their classification of classical reversible circuits.

in the various classes. In fact, every class is exponentially smaller than any class strictly containing it. See Appendix A for details.

(6) **Sporadic Gates.** The process of classification unearthed certain strange classes, which arise from the sometimes complicated interaction of the various invariants which might have otherwise gone overlooked. For example, there are four classes containing the T_4 gate which *require* a generator over at least 4 qubits. Nevertheless, the smallest generator for the class $\langle \mathrm{T}_4, \Gamma, \mathcal{P} \rangle$ is a gate over three qubits. We investigate such a gate in Appendix C. Interestingly, there are no affine gate sets in the classification of classical reversible gates which admit a generator over three bits and no smaller.

1.2 Proof Outline

We can break the proof into a few major steps. First, we introduce the notion of a stabilizer tableau, a binary matrix representation of the function computed by a circuit. We then "guess" all the classes in the classification and designate them by their generators. An examination of the tableaux of the gates in these classes reveals candidate invariants based on the elements of those tableaux. We then prove that these candidate invariants are indeed invariant under the rules of our model. That is to say, if we have two gates whose tableaux satisfy the invariant, then the tableau of their composition satisfies the invariant, and so on for all the other ways to build circuits from gates—tensor products, ancillas, swapping.

At this point, we will have shown that each class has a corresponding invariant, from which the distinctness of the classes in the lattice is immediate. That is, for every pair of distinct classes there is a generator of one that fails to satisfy the invariant of the other. Next, we will show that this correspondence is bijective; that is, for each tableau invariant there is a corresponding class whose generators generate all gates satisfying that tableau invariant. This will have the side effect of producing canonical forms for most of the classes in the lattice.

The challenge remains then to show that our list of classes is exhaustive. Suppose we are given some gate set G. Clearly, the class generated by G is contained in some class in the lattice. Let $\mathcal{C}$ be the smallest class containing all the gates in G. The hope is that we can use gates from G to give explicit circuits for the generators of $\mathcal{C}$. For each class $\mathcal{S} \subset \mathcal{C}$, there must be some gate $g \in G$ which is not in $\mathcal{S}$, otherwise $\mathcal{S}$ would be a smaller class containing G.

The next step of the proof is to focus on some particular gate g, which has failed to satisfy the invariant of S. The key now is the fact that from g we can extract a small generator that *also* fails to satisfy the invariant of S. This is accomplished via the "universal construction" of the gate g, which is a particular circuit built from g and SWAP gates from which these small generators can easily be extracted. Finally, we show that these generators, one for each class $\mathcal{S}$ suffice to generate the class $\mathcal{C}$ itself.

2 Stabilizer Formalism

The one-qubit unitary operations

$$X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \qquad Y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix} \qquad Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$$

are known as *Pauli matrices*. The Pauli matrices are all involutions ($X^2 = Y^2 = Z^2 = I$), and have the following relations between them

$$XY = iZ \qquad YZ = iX \qquad ZX = iY$$
$$YX = -iZ \qquad ZY = -iX \qquad XZ = -iY.$$

It follows that the Pauli matrices generate a discrete group (under multiplication), called the *Pauli group* $\mathcal{P}$, which consists of sixteen elements: $\{I, X, Y, Z\}$ with phases ± 1 or $\pm i$. Then the *Pauli group on n qubits*, $\mathcal{P}_n$, is the set of all n-qubit tensor products of elements from $\mathcal{P}$. We define a *Pauli string* as any element of $\mathcal{P}_n$ with positive phase (i.e., a tensor product of the matrices I, X, Y, Z). We frequently omit the tensor product symbol from Pauli strings and write, e.g., $P_1 \cdots P_n$ where we mean $P_1 \otimes \cdots \otimes P_n$.

The *Clifford group on n qubits*, $\mathcal{C}_n$, is the set of unitary operations which *normalize* $\mathcal{P}_n$ in the group-theoretic sense. That is, $U \in \mathcal{C}_n$ if $UpU^\dagger \in \mathcal{P}_n$ for all $p \in \mathcal{P}_n$. We leave it as a simple exercise to check that $\mathcal{C}_n$ is indeed a group.

A *stabilizer gate* is any unitary in $\bigcup_{n \geq 1} \mathcal{C}_n$. A *stabilizer circuit* is a quantum circuit of stabilizer gates implementing a unitary transformation on some set of qubits, designated the *input/output qubits*, while preserving the state of the remaining *ancilla qubits*. We say that a state $|\psi\rangle$ is *stabilized* by an operation U iff $U|\psi\rangle = |\psi\rangle$. In other words, $|\psi\rangle$ is in the $+1$ eigenspace of U. The Pauli elements and their corresponding stabilized states are below:

$$\begin{array}{ll} X : |+\rangle = \frac{|0\rangle+|1\rangle}{\sqrt{2}} & -X : |-\rangle = \frac{|0\rangle-|1\rangle}{\sqrt{2}} \\ Y : |i\rangle = \frac{|0\rangle+i|1\rangle}{\sqrt{2}} & -Y : |-i\rangle = \frac{|0\rangle-i|1\rangle}{\sqrt{2}} \\ Z : |0\rangle & -Z : |1\rangle \end{array}$$

We call the vectors stabilized by non-identity Pauli elements P and $-P$ the *P-basis*. A *stabilizer state* is any state $U|0\ldots0\rangle$ where U is a stabilizer circuit. For example, $|0\rangle$, $|1\rangle$, $|+\rangle$, $|-\rangle$, $|i\rangle$, and $|-i\rangle$ are the 6 stabilizer states on one qubit. Multi-qubit stabilizer states include $\frac{1}{n}(|0^n\rangle + |1^n\rangle)$ and $\sum_{x \in \{0,1\}^n} |x\rangle$. In general, stabilizer states are of the form (unnormalized) $\sum_{x \in A} (-1)^{q(x)} i^{\ell(x)} |x\rangle$ where A is an affine space over $\mathbb{F}_2$, $q(x)$ is a quadratic form, and $\ell(x)$ is a linear form [11, 17].

3 Gates

Let us introduce some common stabilizer gates. We start with the single-qubit stabilizer gates, which happen to correspond to symmetries of cube (see Figure 1) or octahedron.

3.1 Single-qubit Gates

Face rotations: The Pauli matrices X, Y, and Z (as gates) correspond to 180° rotations about the X, Y, and Z axes respectively. Similarly, we define R_X, R_Y, and R_Z to be 90° rotations (in the counterclockwise direction) about their respective axes. Formally,

$$\mathrm{R}_X = \frac{I - iX}{\sqrt{2}}, \qquad \mathrm{R}_Y = \frac{I - iY}{\sqrt{2}}, \qquad \mathrm{R}_Z = \frac{I - iZ}{\sqrt{2}},$$

although in the case of R_Z (also known as the *phase gate* and often denoted by S or P), a different choice of phase is more conventional. The clockwise rotations are then $\mathrm{R}_X^\dagger$, $\mathrm{R}_Y^\dagger$ and $\mathrm{R}_Z^\dagger$.

Edge rotations: Another symmetry of the cube is to rotate one of the edges 180°. Opposing edges produce the same rotation, so we have six gates: θ_{X+Y}, θ_{X-Y}, θ_{X+Z}, θ_{X-Z}, θ_{Y+Z}, θ_{Y-Z}. We define

$$\theta_{P+Q} = \frac{P+Q}{\sqrt{2}}, \qquad \theta_{P-Q} = \frac{P-Q}{\sqrt{2}},$$

for all Pauli matrices $P \neq Q$. Note that θ_{X+Z} is the well-known *Hadamard gate*, usually denoted by H.

Vertex rotations: The final symmetry is a 120° counterclockwise rotation around one of the diagonals passing through opposite vertices of the cube. The cube has eight vertices, $(\pm 1, \pm 1, \pm 1)$, and we denote the corresponding single-qubit gates Γ_{+++}, Γ_{++-}, ..., Γ_{---}. Algebraically, we define

$$\begin{aligned}
\Gamma_{+++} &= \frac{I - iX - iY - iZ}{2}, \\
\Gamma_{++-} &= \frac{I - iX - iY + iZ}{2}, \\
&\vdots \\
\Gamma_{---} &= \frac{I + iX + iY + iZ}{2}.
\end{aligned}$$

We also define Γ (without subscripts) to be the first gate, Γ_{+++}, since it is the most convenient; conjugation by Γ maps X to Y, Y to Z, and Z to X.

3.2 Multi-qubit Gates

We now introduce the multi-qubit stabilizer gates relevant to the classification. The *SWAP gate*, for instance, which exchanges two qubits. A more interesting example is the *controlled-NOT* or *CNOT gate*, and the *generalized CNOT gates*.

A *generalized CNOT gate* is a two-qubit stabilizer gate of the form

$$C(P,Q) := \frac{I \otimes I + P \otimes I + I \otimes Q - P \otimes Q}{2},$$

where P and Q are Pauli matrices. If the first qubit is in the $+1$ eigenspace of P then $C(P,Q)$ does nothing, but if it is in the -1 eigenspace of P then $C(P,Q)$ applies Q to the second qubit. Of course, the definition is completely symmetric, so you can also view it as applying P to the first qubit when the second qubit is in the -1 eigenspace of Q.

Observe that $C(Z,X)$ is actually the CNOT gate; it applies a NOT gate to the second qubit when the first qubit is $|1\rangle$ and does nothing when the first qubit is $|0\rangle$. Figure 2 shows this equivalence, and illustrates our circuit diagram notation for generalized CNOT gates. Also note that $C(X,Z)$ is a CNOT, but with the opposite orientation (i.e., the second bit controls the first). The rest of the *heterogeneous generalized CNOT gates* (i.e., $C(P,Q)$ where $P \neq Q$) are the natural equivalents of CNOT in different bases.

Similarly, $C(Z,Z)$ is better known as the *controlled-sign gate* (denoted CSIGN), which flips the sign on input $|11\rangle$, but does nothing otherwise. The *homogeneous generalized CNOT gates*

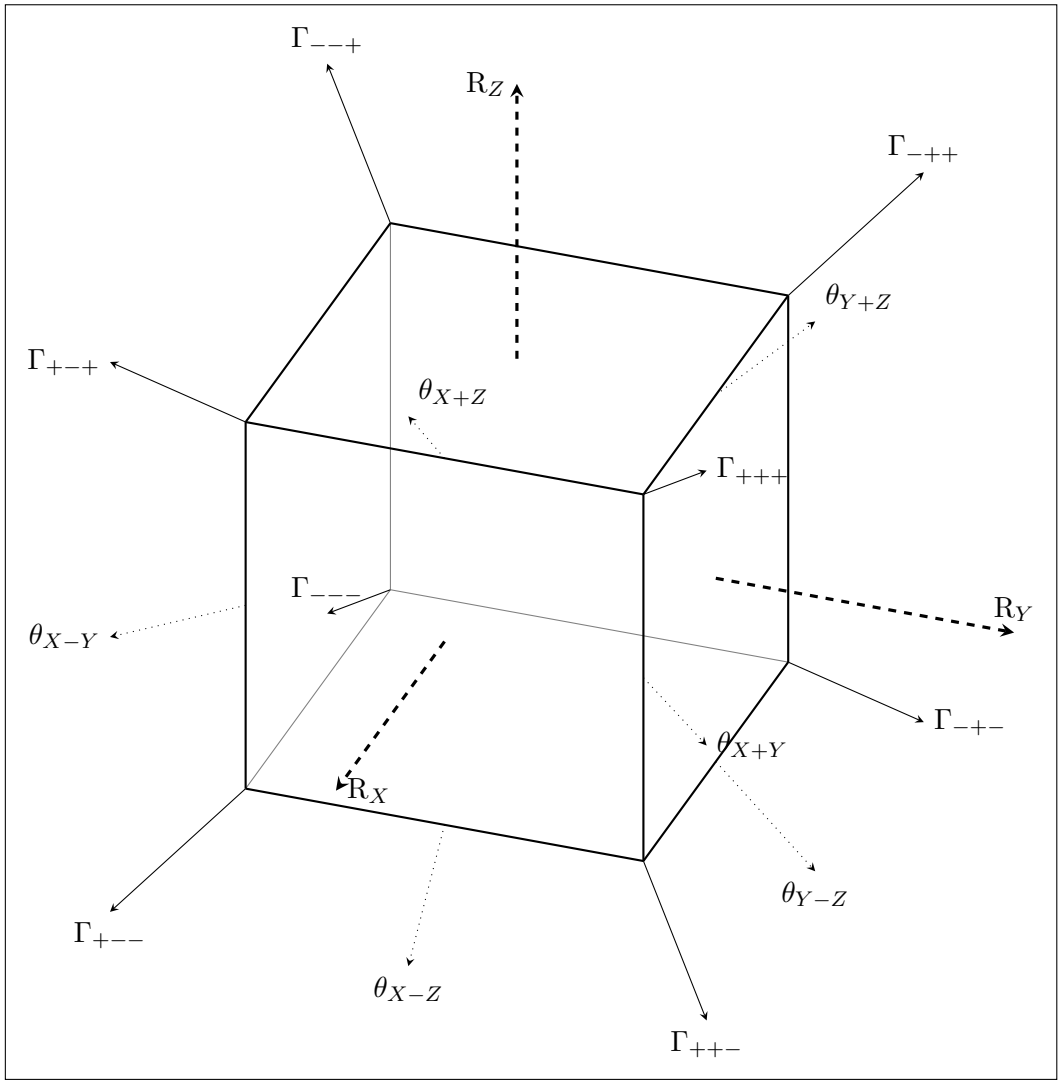

Figure 1: Single-qubit gates as symmetries of the cube

Gate	Tableau	Unitary Matrix
X	$\left(\begin{array}{cc\|c} 1 & 0 & 0 \\ 0 & 1 & 1 \end{array}\right)$	$\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$
Y	$\left(\begin{array}{cc\|c} 1 & 0 & 1 \\ 0 & 1 & 1 \end{array}\right)$	$\begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}$
Z	$\left(\begin{array}{cc\|c} 1 & 0 & 1 \\ 0 & 1 & 0 \end{array}\right)$	$\begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$
R_X	$\left(\begin{array}{cc\|c} 1 & 0 & 0 \\ 1 & 1 & 1 \end{array}\right)$	$\frac{1}{\sqrt{2}}\begin{pmatrix} 1 & -i \\ -i & 1 \end{pmatrix}$
R_Y	$\left(\begin{array}{cc\|c} 0 & 1 & 1 \\ 1 & 0 & 0 \end{array}\right)$	$\frac{1}{\sqrt{2}}\begin{pmatrix} 1 & -1 \\ 1 & 1 \end{pmatrix}$
$\mathrm{R}_Z = S$	$\left(\begin{array}{cc\|c} 1 & 1 & 0 \\ 0 & 1 & 0 \end{array}\right)$	$\frac{1-i}{\sqrt{2}}\begin{pmatrix} 1 & 0 \\ 0 & i \end{pmatrix}$
$\theta_{XZ} = \theta_{X+Z} = H$	$\left(\begin{array}{cc\|c} 0 & 1 & 0 \\ 1 & 0 & 0 \end{array}\right)$	$\frac{1}{\sqrt{2}}\begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}$
θ_{X-Z}	$\left(\begin{array}{cc\|c} 0 & 1 & 1 \\ 1 & 0 & 1 \end{array}\right)$	$\frac{1}{\sqrt{2}}\begin{pmatrix} -1 & 1 \\ 1 & 1 \end{pmatrix}$
$\Gamma_{+++} = \Gamma$	$\left(\begin{array}{cc\|c} 1 & 1 & 0 \\ 1 & 0 & 0 \end{array}\right)$	$\frac{1-i}{2}\begin{pmatrix} 1 & -i \\ 1 & i \end{pmatrix}$

Table 1: Single-qubit gates

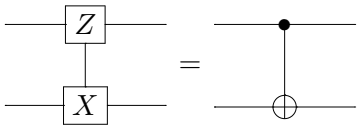

Figure 2: CNOT expressed as a $C(Z, X)$ gate

(i.e., $C(P, P)$ for some P) are quite different from heterogeneous CNOT gates. For instance, when one CNOT targets the control qubit of another CNOT then it matters which gate is applied first. On the other hand, two CSIGN gates will always commute, whether or not they have a qubit in common.

It turns out that every two-qubit stabilizer gate is equivalent to some combination of one qubit stabilizer gates and generalized CNOT gates (see Appendix D). Most classes of stabilizer gates are of this form, but there are a handful of classes which require larger generators such as the T_4 gate [3, 14]. For all $k \geq 1$, we define T_{2k} as the $2k$-bit classical gate that flips all bits of the input when the parity of the input is odd and does nothing when the parity of the input is even. In particular, T_2 is the lowly SWAP gate. Notice that T_{2k} is an orthogonal linear function of the input bits. The

Gate	Tableau	Unitary Matrix on computational basis
$C(X,X)$	$\left(\begin{array}{cc\|cc} 1&0&0&0\\ 0&1&1&0\\ \hline 0&0&1&0\\ 1&0&0&1 \end{array}\right)$	$\frac{1}{2}\begin{pmatrix} 1&1&1&-1\\ 1&1&-1&1\\ 1&-1&1&1\\ -1&1&1&1 \end{pmatrix}$
$C(Y,Y)$	$\left(\begin{array}{cc\|cc} 1&0&1&1\\ 0&1&1&1\\ \hline 1&1&1&0\\ 1&1&0&1 \end{array}\right)$	$\frac{1}{2}\begin{pmatrix} 1&-i&-i&1\\ i&1&-1&-i\\ i&-1&1&-i\\ 1&i&i&1 \end{pmatrix}$
$C(Z,Z) = \text{CSIGN}$	$\left(\begin{array}{cc\|cc} 1&0&0&1\\ 0&1&0&0\\ \hline 0&1&1&0\\ 0&0&0&1 \end{array}\right)$	$\begin{pmatrix} 1&0&0&0\\ 0&1&0&0\\ 0&0&1&0\\ 0&0&0&-1 \end{pmatrix}$
$C(Y,X)$	$\left(\begin{array}{cc\|cc} 1&0&1&0\\ 0&1&1&0\\ \hline 0&0&1&0\\ 1&1&0&1 \end{array}\right)$	$\frac{1}{2}\begin{pmatrix} 1&1&-i&i\\ 1&1&i&-i\\ i&-i&1&1\\ -i&i&1&1 \end{pmatrix}$
$C(Z,Y)$	$\left(\begin{array}{cc\|cc} 1&0&1&1\\ 0&1&0&0\\ \hline 0&1&1&0\\ 0&1&0&1 \end{array}\right)$	$\begin{pmatrix} 1&0&0&0\\ 0&1&0&0\\ 0&0&0&-i\\ 0&0&i&0 \end{pmatrix}$
$C(X,Z) = \text{CNOT}$	$\left(\begin{array}{cc\|cc} 1&0&0&0\\ 0&1&0&1\\ \hline 1&0&1&0\\ 0&0&0&1 \end{array}\right)$	$\begin{pmatrix} 1&0&0&0\\ 0&0&0&1\\ 0&0&1&0\\ 0&1&0&0 \end{pmatrix}$

Table 2: Two qubit gates. The sign bits are all 0 in the above tableaux so they are omitted.

matrix of the T_4 gate (over $\mathbb{F}_2$) is

$$\begin{pmatrix} 0&1&1&1\\ 1&0&1&1\\ 1&1&0&1\\ 1&1&1&0 \end{pmatrix}.$$

In the quantum setting, let the T_{2k} gate simply apply this transformation to the Z basis vectors.

4 Tableaux

Observe that the matrices I, X, Y, Z are linearly independent, and therefore form a basis for the 2×2 complex matrices. It follows that $\mathcal{P}_n$ spans all $2^n \times 2^n$ complex matrices. Hence, any unitary operation on n qubits can be characterized by its action on $\mathcal{P}_n$. In particular, any gate is characterized by how it acts on

$$p_1 = XI\cdots I, p_2 = ZI\cdots I, \ldots, p_{2n-1} = I\cdots IX, p_{2n} = I\cdots IZ.$$

We call this list the *Pauli basis on n qubits*, since one can write any element of $\mathcal{P}_n$ as a product of basis elements times a phase (± 1 or $\pm i$).

Now suppose we are given a stabilizer gate, $U \in \mathcal{C}_n$. By definition, stabilizer gates will map each Pauli basis element to something in $\mathcal{P}_n$, which can be written as a product of basis elements times a phase. That is,

$$U p_j U^\dagger = \alpha_j \prod_{k=1}^{2n} p_k^{M_{jk}}$$

for some bits $M_{j1}, \ldots, M_{j(2n)} \in \{0,1\}$ and some phase[4] $\alpha_j \in \{\pm 1, \pm i\}$. The *tableau* for U is a succinct representation for U consisting of the binary matrix $M = [M_{jk}]$, and some representation of the phases $\alpha_1, \ldots, \alpha_{2n}$.

It turns out that U maps p_j (or any Pauli string) to ± 1 times a Pauli string, never $\pm i$ times a Pauli string. This follows from the fact that the square of any Pauli string is $I \cdots I$, so

$$(U p_j U^\dagger)^2 = U p_j^2 U^\dagger = U(I \cdots I) U^\dagger = +I \cdots I.$$

If the phase in front of the Pauli string in $U p_j U^\dagger$ were $\pm i$, then squaring it would produce a negative phase. Therefore, $U p_j U^\dagger$ is ± 1 times a Pauli string.

Unfortunately, α_j may still be any one of $\{\pm 1, \pm i\}$. This is because the product of $p_{2k-1} p_{2k}$ is $I \cdots I(XZ) I \cdots I = -iI \cdots IYI \cdots I$, with an awkward $-i$ phase. Once we cancel the extra factors of i from α_j, we are left with

$$(-1)^{v_j} := \alpha_j \prod_{k=1}^{n} (-i)^{M_{j(2k-1)} M_{j(2k)}},$$

where $v_j \in \{0,1\}$ is the *phase bit for row j*. For example, if $U p_1 U^\dagger$ is $YI \cdots I$ then we have $M_{11} = M_{12} = 1$ and $v_1 = 0$. Then the complete tableau for U is the matrix $M = [M_{jk}]$ and the vector of bits $v = [v_j]$, which we typically write as

$$\left(\begin{array}{ccc|c} M_{11} & \cdots & M_{1(2n)} & v_1 \\ \vdots & \ddots & \vdots & \vdots \\ M_{(2n)1} & \cdots & M_{(2n)(2n)} & v_{2n} \end{array}\right).$$

Our ordering of the basis elements (which differs from other presentations) puts Pauli strings on the same qubit (e.g., $XI \cdots I$ and $ZI \cdots I$) side-by-side in the matrix, so the 2×2 submatrix

$$\begin{pmatrix} M_{2i-1,2j-1} & M_{2i-1,2j} \\ M_{2i,2j-1} & M_{2i,2j} \end{pmatrix}$$

completely describes how the ith qubit of the input affects the jth qubit of the output. In fact, it will be fruitful to think of the tableau as an $n \times n$ matrix of 2×2 blocks, along with a vector of $2n$ *phase bits*. To be clear, the blocks come from $\mathsf{R} := \mathbb{F}_2^{2\times 2}$, the ring of 2×2 matrices over the field of two elements, $\mathbb{F}_2$. Then the tableau is a matrix in $\mathsf{R}^{n \times n}$ (the $n \times n$ matrices over the ring R), combined with a vector of phase bits in $\mathbb{F}_2^{2n}$. Each row of the matrix is associated with a pair of phase bits from the vector.

[4]Note that the order of the terms in the product matters, since the Pauli basis elements do not necessarily commute, so we assume the terms are in the natural order from $p_1^{M_{j1}}$ up to $p_{2n}^{M_{j(2n)}}$.

However, not every matrix in $\mathsf{R}^{n\times n}$ corresponds to a stabilizer circuit due to unitarity constraints. To best express these constraints, we define a unary operation * on R such that

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix}^* = \begin{pmatrix} d & b \\ c & a \end{pmatrix}.$$

The * operator has the property that

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix}\begin{pmatrix} a & b \\ c & d \end{pmatrix}^* = \begin{pmatrix} a & b \\ c & d \end{pmatrix}\begin{pmatrix} d & b \\ c & a \end{pmatrix} = \begin{pmatrix} ad+bc & 0 \\ 0 & ad+bc \end{pmatrix} = I\begin{vmatrix} a & b \\ c & d \end{vmatrix}.$$

Additionally,

$$\begin{aligned} I^* &= I, \\ (M+N)^* &= M^* + N^*, \\ (MN)^* &= N^*M^*, \\ (M^*)^* &= M, \end{aligned}$$

so * makes R a **-ring* or *ring with involution*. We also extend * to an operation on matrices (over R) which applies * to each entry and then transposes the matrix. It turns out that a tableau represents a unitary operation if and only if the matrix $M \in \mathsf{R}^{n\times n}$ satisfies $MM^* = M^*M = I$. This (intentionally) resembles the definition of a unitary matrix ($UU^\dagger = U^\dagger U = I$), we will call this the *unitarity condition*, and it corresponds to the unitarity of U as a gate, but M is certainly not a traditional complex unitary matrix (nor unitary over some finite field with conjugation).

4.1 Correspondence between Gates and Tableaux

We will find it useful to switch between gates and tableaux, as one notion often informs the other. In light of that, define $\mathcal{T}(g)$ to be the tableau of a stabilizer gate g. In fact, we will use $\mathcal{M}(g)$ to denote only the matrix part of the tableau of g, as the phase bits of the tableau often prove to be irrelevant. Indeed, most non-degenerate gate sets generate the Pauli group, which alone suffices to set the phase bits of the tableau arbitrarily as follows: applying X to qubit j negates v_{2j} and applying Z to qubit j negates v_{2j-1}. Furthermore, there is a surprising connection between individual entries of tableaux and elementary stabilizer operations that can be extracted from them (see Section 9).

If $a \in \mathsf{R}$ is invertible, then let $\mathcal{G}(a)$ be the single-qubit gate with $\mathcal{M}(\mathcal{G}(a)) = a$ and zeros for phase bits. These are the gates in the first row of Table 3. Let $\mathcal{G}(a,i)$ be the gate $\mathcal{G}(a)$ applied to the ith qubit.

If $b \in \mathsf{R}$ is noninvertible, define $\mathrm{CNOT}(b,i,j)$ to be the generalized CNOT on qubits i and j corresponding to the singular matrix b. These are the gates in the first row of Table 4. The tableau $\mathcal{T}(\mathrm{CNOT}(b,i,j))$ is the identity tableau except for b^* and b in positions (i,j) and (j,i), respectively. We use the circuit in Figure 3 to designate such a gate.

Finally, we would like to have a direct way to compose two circuits by a simple operation on their tableaux. Suppose we wish to compute the composition of circuits C_1 and C_2. To compute $\mathcal{T}(C_2 \circ C_1)$, we must compute, for each Pauli basis element p_j, the product $C_2C_1p_jC_1^\dagger C_2^\dagger$. First consider the jth row of $\mathcal{T}(C_1)$, which gives

$$C_1 p_j C_1^\dagger = \alpha_j \prod_{k=1}^{2n} p_k^{M^{(1)}_{jk}},$$

	$\left(\begin{smallmatrix}1&0\\0&1\end{smallmatrix}\right)$	$\left(\begin{smallmatrix}1&0\\1&1\end{smallmatrix}\right)$	$\left(\begin{smallmatrix}0&1\\1&0\end{smallmatrix}\right)$	$\left(\begin{smallmatrix}1&1\\0&1\end{smallmatrix}\right)$	$\left(\begin{smallmatrix}1&1\\1&0\end{smallmatrix}\right)$	$\left(\begin{smallmatrix}0&1\\1&1\end{smallmatrix}\right)$
$\left(\begin{smallmatrix}0\\0\end{smallmatrix}\right)$	I	$\mathrm{R}_X^\dagger$	θ_{X+Z}	R_Z	Γ_{+++}	Γ_{---}
$\left(\begin{smallmatrix}0\\1\end{smallmatrix}\right)$	X	R_X	$\mathrm{R}_Y^\dagger$	θ_{X+Y}	Γ_{--+}	Γ_{+-+}
$\left(\begin{smallmatrix}1\\0\end{smallmatrix}\right)$	Z	θ_{Y+Z}	R_Y	$\mathrm{R}_Z^\dagger$	Γ_{-+-}	Γ_{-++}
$\left(\begin{smallmatrix}1\\1\end{smallmatrix}\right)$	Y	θ_{Y-Z}	θ_{X-Z}	θ_{X-Y}	Γ_{+--}	Γ_{++-}

Table 3: Invertible tableau elements and the corresponding single-qubit gates produced by the universal construction. Row of the table corresponds to the sign bit of the row of the tableau in which the element occurs.

Element	$\left(\begin{smallmatrix}0&0\\1&0\end{smallmatrix}\right)$	$\left(\begin{smallmatrix}1&1\\1&1\end{smallmatrix}\right)$	$\left(\begin{smallmatrix}0&1\\0&0\end{smallmatrix}\right)$	$\left(\begin{smallmatrix}1&0\\0&0\end{smallmatrix}\right)$ / $\left(\begin{smallmatrix}0&0\\0&1\end{smallmatrix}\right)$	$\left(\begin{smallmatrix}1&0\\1&0\end{smallmatrix}\right)$ / $\left(\begin{smallmatrix}0&0\\1&1\end{smallmatrix}\right)$	$\left(\begin{smallmatrix}1&1\\0&0\end{smallmatrix}\right)$ / $\left(\begin{smallmatrix}0&1\\0&1\end{smallmatrix}\right)$
Gen. CNOT	$C(X,X)$	$C(Y,Y)$	$C(Z,Z)$	$C(X,Z)$	$C(Y,X)$	$C(Z,Y)$

Table 4: Noninvertible tableau elements and the corresponding generalized CNOT gates produced by the universal construction.

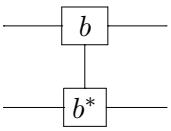

Figure 3: Circuit diagram for CNOT$(b,1,2)$ gate.

where $M^{(1)}$ is the binary representation of $\mathcal{M}(C_1)$, and α_j is the phase. Similarly,

$$C_2 p_j C_2^\dagger = \beta_j \prod_{k=1}^{2n} p_k^{M^{(2)}_{jk}}.$$

Therefore,

$$\begin{aligned} C_2 C_1 p_j C_1^\dagger C_2^\dagger &= C_2 \left(\alpha_j \prod_{k=1}^{2n} p_k^{M^{(1)}_{jk}}\right) C_2^\dagger = \alpha_j \prod_{k=1}^{2n} \left(C_2 p_k^{M^{(1)}_{jk}} C_2^\dagger\right) = \alpha_j \prod_{k=1}^{2n} \left(\beta_k \prod_{\ell=1}^{2n} p_\ell^{M^{(2)}_{k\ell}}\right)^{M^{(1)}_{jk}} \\ &= \alpha_j \prod_{k=1}^{2n} \left(\beta_k^{M^{(1)}_{jk}} \prod_{\ell=1}^{2n} p_\ell^{M^{(1)}_{jk} M^{(2)}_{k\ell}}\right) \propto \prod_{\ell=1}^{2n}\prod_{k=1}^{2n} p_\ell^{M^{(1)}_{jk} M^{(2)}_{k\ell}} \propto \prod_{\ell=1}^{2n} p_\ell^{\bigoplus_{k=1}^{2n} M^{(1)}_{jk} M^{(2)}_{k\ell}} \\ &= \prod_{\ell=1}^{2n} p_\ell^{[M^{(1)}M^{(2)}]_{j\ell}} \end{aligned}$$

Notice that this implies $\mathcal{M}(C_2 \circ C_1) = \mathcal{M}(C_1)\mathcal{M}(C_2)$. Since it is cumbersome to write out explicitly, we did not include the exact phases in the above calculation. Nevertheless, one can compute the phase bits by tracking the intermediate steps in the above calculation, which includes the multiplication of Pauli basis elements.

5 Classes

Informally, the class generated by a set of stabilizer gates is the collection of stabilizer circuits which can be performed with those gates. Our goal is to determine the set of classes and the relationships between them. We now must lay out the precise operations we use to build circuits. Composition and tensor product are obvious, but let us introduce and justify our rules concerning the reordering of qubits and ancillary workspace.

First, we will always allow the swapping of qubits. This allows us to consider gates without needing to specify the qubits on which they must be applied. Indeed, we can relabel the input wires however we like. Secondly, we allow the use of quantum ancillas, that is, arbitrary states that we can "hardwire" into the computation. We only stipulate that these ancilla qubits be returned to their initial configuration at the end of the computation.

These ancilla inputs can be viewed as the workspace of the computation. If we allow no ancilla inputs at all then, of course, a gate set consisting of gates over many inputs can't possibly hope to simulate a gate over a smaller number of qubits. Furthermore, if we want to apply a stabilizer function as a subroutine of a general quantum computation (as we would, say, for a fault-tolerant scheme), then we need that the ancilla states do not depend on the input. Otherwise, they would destroy the quantum coherence of the computation.

Finally, we turn to the question of the ancillas themselves. The weakest assumption one could make is that the ancillas are initialized to an unknown state, which the circuit may only change temporarily. This is somewhat artificial since we can, at the very least, initialize the workspace to the all-zeros state. Other classifications suggest that without this assumption the problem becomes *dramatically* more difficult (for example, Post's infinite lattice collapses to a small finite one [3]). A slightly stronger assumption would be to allow ancillas initialized to computational basis states, but this would break symmetry by introducing a bias towards the Z-basis.[5]

A next natural step would be to permit ancillas initialized to arbitrary stabilizer states. Although this would appear to be circular (i.e., stabilizer gates are necessary to implement stabilizer states, which we then use as ancillas in stabilizer circuits), the reusability of our ancillas implies that even if the states are difficult to construct, at least we only have to construct them once. Unfortunately, we are unable to complete the classification in its entirety under this ancilla model. However, we have reduced the problem to finding a stabilizer state which is stabilized by Γ and a permutation, and we conjecture that such a state exists (see Section 11). Moreover, the conjectured classification matches the one we will present (under a stronger ancilla model).

Finally, we arrive at our chosen model; that is, ancillas initialized to arbitrary quantum states. A priori, these states could be arbitrarily large, and arbitrarily complicated to construct, which is clearly undesirable. It turns out, however, the classification only requires finitely many one or two qubit states, in particular, the eigenstates of the single-qubit stabilizer gates and states that are locally equivalent to the Bell state. For comparison, we have determined the reversible gate lattice under quantum ancillas in Appendix B, and observe that, in some cases, arbitrarily large, entangled states are necessary.

It is worth noting that there is a long line of work showing that weak gate sets, including stabilizer gates, are universal for quantum computation when given access to magic states [8, 18, 20]. Importantly, magic states do not need to be preserved after the computation, contrary to our model.

[5] We can fix the bias by allowing all single-qubit ancillary states: $|0\rangle, |1\rangle, |+\rangle, |-1\rangle, |i\rangle, |-i\rangle$. This introduces new classes such as $\langle \theta_{X+Z} \otimes \theta_{X+Z} \rangle$, but we leave the classification under these assumptions as an open problem.

Indeed, our classification shows that arbitrary quantum states cannot boost the power of stabilizer gates beyond the Clifford group.

One can increase the power of ancillas by letting them change over the course of the computation, as long as the change is independent of the input (that is, from some constant initial state to some possibly different constant final state). Indeed, in the classification of reversible gates [3], these "loose ancillas" collapse a few pairs of classes. Nevertheless, we show that even with loose ancillas, our classification (as presented) still holds. When we formally define class invariants in Section 6 (see Theorem 3 in particular), we will see that the invariants hold under the loose quantum ancilla model, and therefore hold for all weaker models.

Now let us formally define a *class* $\mathcal{C}$ to be a set of stabilizer gates satisfying the following four rules:

(1.) Composition Rule $\mathcal{C}$ is closed under composition of gates.

(2.) Tensor Rule $\mathcal{C}$ is closed under tensor product of gates.

(3.) Swap Rule $\mathcal{C}$ contains the SWAP gate.

(4.) Ancilla Rule $\mathcal{C}$ is closed under ancillas. More precisely, if we have a gate $g \in \mathcal{C}$ and state $|\psi\rangle$, and there exists g' such that $g(|x\rangle \otimes |\psi\rangle) = g'(|x\rangle) \otimes |\psi\rangle$ for all inputs $|x\rangle$ (up to a global phase), then g' is a member of $\mathcal{C}$.

Let's first see a few simple consequences of the model.

Proposition 1. *Let $\mathcal{C}$ be a class of stabilizer gates. Then $\mathcal{C}$ contains the n qubit identity gate for any n.*

Proof. First, $\mathcal{C}$ contains SWAP. It follows that $\mathcal{C}$ contains the two qubit identity gate since it is the composition SWAP $\circ$ SWAP. By the ancilla rule, we can remove a qubit from the two qubit identity using *any* one-qubit state. Hence, the one qubit identity gate is in $\mathcal{C}$. Finally, $\mathcal{C}$ must contain the n-qubit identity gate because it is the tensor product of n one-qubit gates. □

Proposition 2. *Let $\mathcal{C}$ be a class of stabilizer gates. For any $g \in \mathcal{C}$, the inverse, g^{-1}, belongs to $\mathcal{C}$.*

Proof. Consider the sequence

$$g, g \circ g, g \circ g \circ g, \ldots, g^n, \ldots.$$

Since there are finitely many stabilizer gates on n qubits (certainly finitely many tableaux, and one gate per tableau), the sequence must eventually repeat. That is, $g^i = g^j$ for some $1 \leq i < j$. Since every stabilizer gate has an inverse, we conclude that $1 = g^0 = g^{j-i}$, and hence $g^{-1} = g^{j-i-1}$. In other words, g^{-1} is a stabilizer gate, and g^{-1} is a (positive) power of g and therefore in $\mathcal{C}$. □

The most practical way to talk about a class $\mathcal{C}$ is by its *generators*. We say a set of gates G *generates* a class $\mathcal{C}$ if $G \subseteq \mathcal{C}$ and every class containing G also contains $\mathcal{C}$. We introduce the notation $\langle \cdot \rangle$ for the class generated by a set of gates. Similarly, we say that G generates a specific gate g if $g \in \langle G \rangle$.

Our goal is therefore to identify all stabilizer gate classes, determine their generators, and diagram the relationships between classes. As it turns out, there are 57 different classes, which we have split across Figure 4 (which contains the classes with single-qubit generators) and Figure 5 (which contains the multi-qubit classes). Each class is labelled by a set of generators for that class,

except for ALL, the class of all stabilizer gates; $\top$, the class of all single-qubit stabilizer gates; and $\bot$, the class generated by the empty set. Additionally, we abbreviate some class names in Figure 4:

- θ_{+++}, θ_{+--}, θ_{-+-}, θ_{--+} denote the single-qubit classes containing Γ_{+++}, Γ_{+--}, Γ_{-+-}, and Γ_{--+} respectively, and three θ gates each, as indicated by the gray lines.
- θ_{xy} abbreviates θ_{x+y} or θ_{x-y} (it contains both) and similarly for θ_{xz} and θ_{yz}.

Some of the lines in Figure 4 are gray and dotted, not for any technical reason, but because the lattice would be unreadable otherwise.

In Figure 5 each class includes the label of the single-qubit subgroup, even when not all of the single-qubit generators are necessary to generate the class. This is intended to make the relationship between the degenerate and non-degenerate lattices clearer. For example, T_4 generates the Pauli group, $\mathcal{P}$, on its own (Lemma 11), but we label the class $\langle \mathrm{T}_4, \mathcal{P} \rangle$ to make it clear that the class $\langle \mathrm{T}_4 \rangle$ is above $\langle \mathcal{P} \rangle$ in the lattice.

6 Invariants

Until now, we have defined each class in terms of the generators for that class. It turns out that each class can also be characterized as the set of all gates satisfying a collection of invariants. Section 7 formalizes this equivalence. This section focuses on the form of the invariants themselves.

Informally, an *invariant* is a property of gates, readily apparent from their tableaux, which is preserved by the circuit building operations. In other words, if a collection of gates all satisfy a particular invariant then any circuit constructed from those gates must also satisfy the invariant. All our invariants are formally defined from the tableaux, but for now we give the following informal descriptions to make the intuition for each invariant clear.

X-, Y-, or Z-preserving: We say a stabilizer gate is *Z-preserving* if it maps Z-basis states to Z-basis states, possibly with a change of phase. The Z-preserving gates include all (classical) reversible gates (e.g., X, CNOT, and T_4), gates which only manipulate the sign (e.g., R_Z and CSIGN), and combinations of the two.

Symmetrically, there are X-preserving gates mapping X-basis states to X-basis states, and Y-preserving gates gates mapping Y-basis states to Y-basis states. Our definitions of classes, gates, invariants, etc., are completely symmetric with respect to X, Y and Z basis, so if some gate or class is X-preserving (resp. Y-preserving or Z-preserving), then there must be a corresponding gate or class which is Y-preserving (resp. Z-preserving or X-preserving). We will often appeal to this symmetry to simplify proofs.

Note that a gate can be any combination of X-, Y-, and Z-preserving. For instance, T_4 is X-, Y-, and Z-preserving; CNOT is X-preserving and Z-preserving but not Y-preserving (similarly $C(X, Y)$ and $C(Y, Z)$ fail to be Z-preserving and X-preserving respectively); R_Z is Z-preserving only (similarly R_X is X-preserving and R_Y is Y-preserving); and Γ is not X-, Y-, or Z-preserving.

Egalitarian We say a gate is *egalitarian* if it is fixed (up to a Pauli operation on each qubit) by the aforementioned $X/Y/Z$ symmetry, that is, the symmetry arising from conjugating all qubits by Γ (which cycles X to Y, Y to Z, and Z to X). In particular, this implies that if

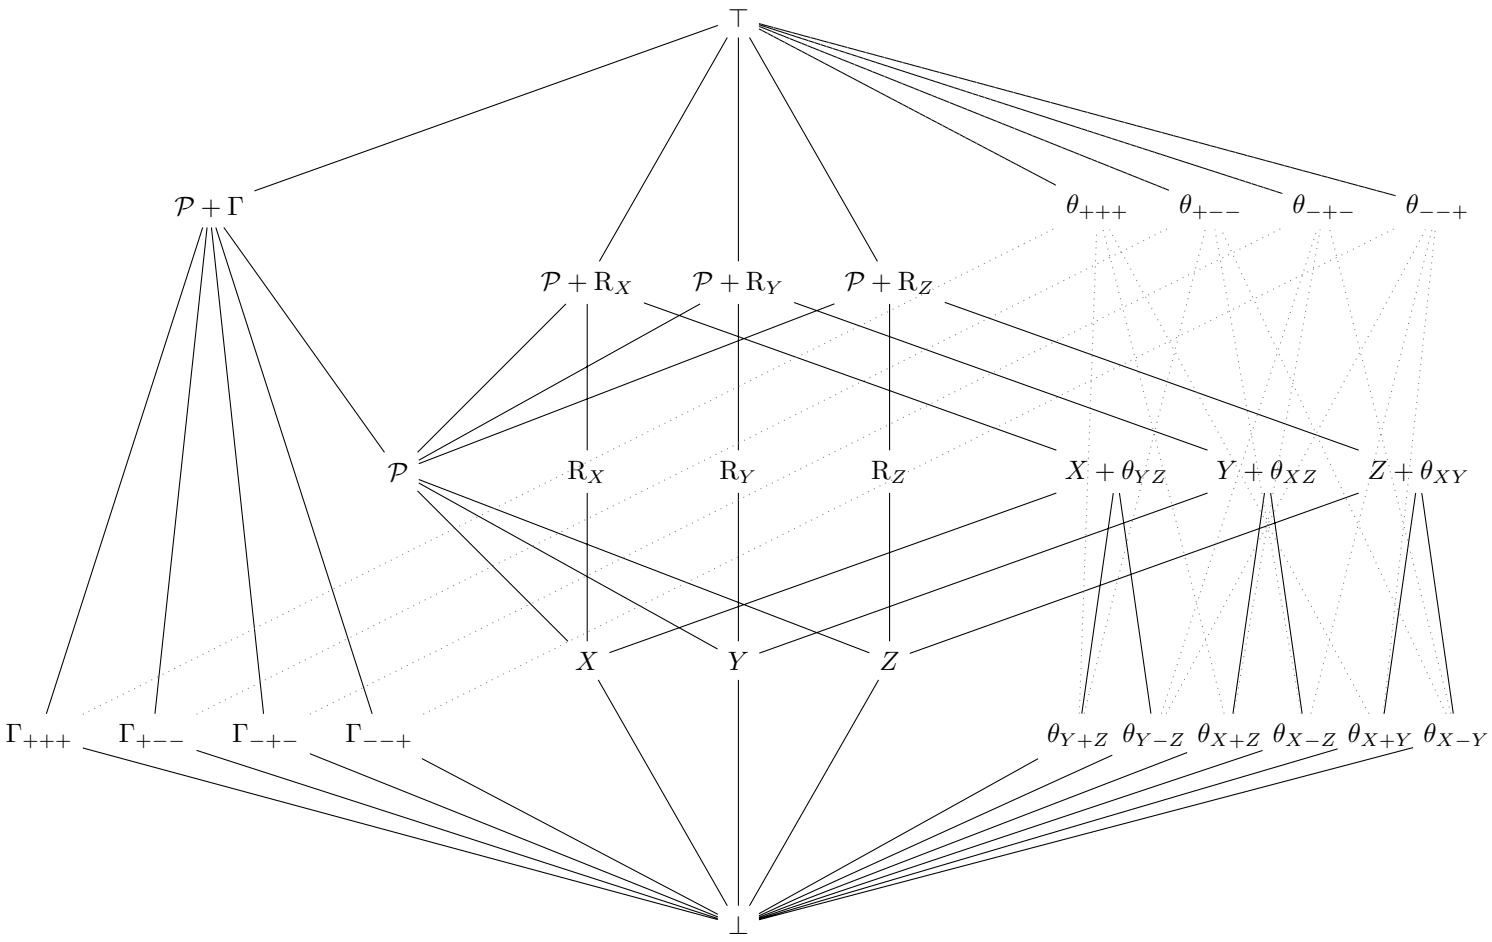

Figure 4: The inclusion lattice of degenerate stabilizer gate classes

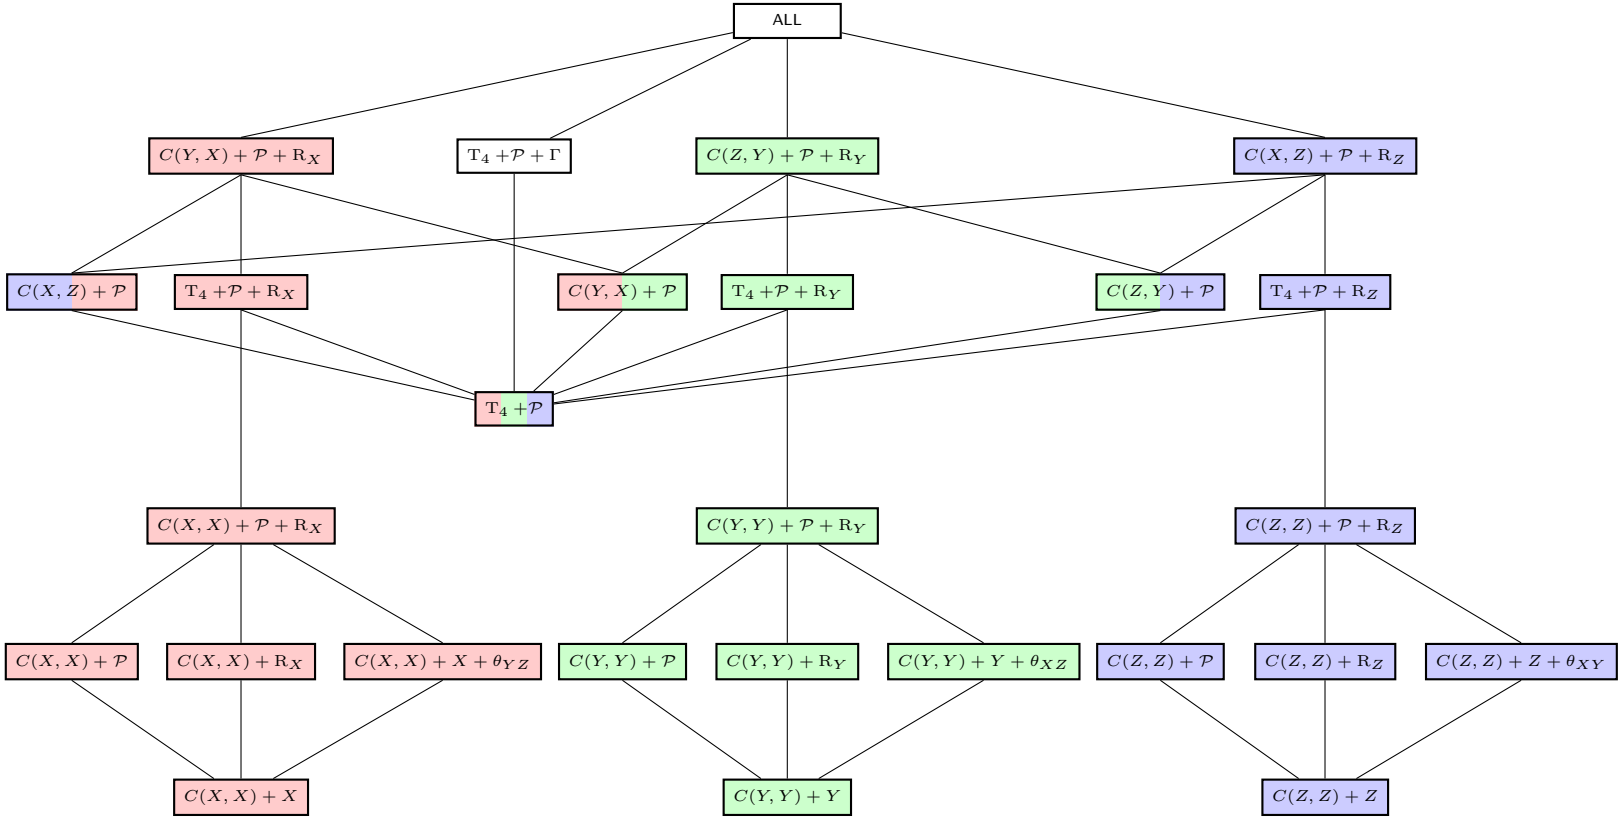

Figure 5: The inclusion lattice of non-degenerate stabilizer gate classes. Red, green, blue denote X-, Y-, and Z-preserving, respectively.

egalitarian operation U maps Pauli string P to $Q = UPU^\dagger$ under conjugation, then U maps $\Gamma P \Gamma^\dagger$ to

$$U\Gamma P\Gamma^\dagger U^\dagger \propto \Gamma U\Gamma^\dagger \Gamma P\Gamma^\dagger \Gamma U^\dagger \Gamma^\dagger = \Gamma UPU^\dagger \Gamma^\dagger = \Gamma Q\Gamma^\dagger.$$

Not only are Γ and the Pauli matrices themselves egalitarian, but so is T_4.

Degenerate: We say a gate is *degenerate* if each input affects only one output. More precisely, when applying the gate to a string of Paulis, changing one Pauli in the input will change exactly one Pauli in the output. All single-qubit gates are degenerate, and all degenerate gates can be composed of single-qubit gates and SWAP gates.

X-, Y-, or Z-degenerate: A gate is *Z-degenerate* if it is Z-preserving and flipping any bit of a classical (Z-basis) input to the gate causes exactly one bit of the output to flip. The gate may or may not affect the phase. This class includes several Z-preserving single qubit gates, like R_Z, the Pauli operations, and θ_{X+Y}. It also includes CSIGN because this gate *only* affects phase, but CNOT is not Z-degenerate because flipping the control bit changes both outputs. Notice that CSIGN is Z-degenerate, but not degenerate. We define X-degenerate and Y-degenerate symmetrically.

X-, Y-, or Z-orthogonal: A gate G is *Z-orthogonal* if it can be built from T_4 and Z-preserving single-qubit gates. The term "orthogonal" comes from the fact that T_4 is an orthogonal linear transformation in the Z-basis, but not all Z-orthogonal gates are literally orthogonal transformations in the Z-basis (see, for example, Lemma 11). Similarly for X-orthogonal and Y-orthogonal.

Single Qubit Gates: There are thirty different classes of single-qubit gates. All of these classes are degenerate, and some can be distinguished by the other invariants above. However, many single-qubit invariants depend on the phase bits of the tableau. For instance, the tableau of θ_{X+Y}, θ_{X-Y}, and R_Z all have the same matrix part, $\left(\begin{smallmatrix} 1 & 1 \\ 0 & 1 \end{smallmatrix}\right)$, but generate three distinct classes. One can write down explicit invariants for these classes where the phase bits are correlated to the tableau entries, but in most cases we present a single-qubit class as a subgroup of the symmetries of the cube/octahedron, as shown in Figure 1.

6.1 Formal invariants

An *invariant* is a property of tableaux which is preserved by the four circuit-building rules.

Swap Rule: Every class contains the SWAP gate, so every invariant we propose must be satisfied by the tableau for SWAP.

Composition Rule: If the invariant holds for two gates, then it must hold for their composition. We have seen that the tableau for the composition of two gates is essentially the matrix product of the two tableau, except for the phase bits (which are significantly more complicated to update).

Tensor Rule: The tensor product of two gates satisfying the invariant must also satisfy the invariant. Note that the tableau of the tensor product is the direct sum of the tableaux, and phase bits are inherited from the sub-tableaux in the natural way.

Ancilla Rule: The invariant must be preserved when some qubits are used as ancillas. It turns out the ancilla operation reduces the tableau to a submatrix (of non-ancilla rows and columns) and under certain conditions, the corresponding subset of the phase bits. This is somewhat technical, so we prove it in Theorem 3 below.

Theorem 3. *Let G be a stabilizer gate on n qubits, and suppose there exist states $|\psi\rangle$ and $|\psi'\rangle$ such that*

$$G(|x\rangle \otimes |\psi\rangle) = H(|x\rangle) \otimes \left|\psi'\right\rangle$$

for all $|x\rangle$, for some unitary H on m-qubits. In particular, this is true if we use the ancilla rule to reduce G to H, where $|\psi\rangle = |\psi'\rangle$ is the ancilla state. Then

1. *H is a stabilizer operation,*
2. *$\mathcal{M}(H)$ is obtained by removing the rows and columns corresponding to the ancilla bits from $\mathcal{M}(G)$,*
3. *If every bit (in $\mathcal{M}(G)$ as a binary matrix) we remove from a row is zero, then the phase bit for that row is the same in $\mathcal{T}(G)$ and $\mathcal{T}(H)$.*

Proof. Let $P \in \mathcal{P}_m$. Then for all $|x\rangle$,

$$G(P|x\rangle \otimes |\psi\rangle) = H(P|x\rangle) \otimes \left|\psi'\right\rangle.$$

On the other hand, G is a stabilizer gate, so conjugating the Pauli string $P \otimes I^{n-m}$ by G produces $Q \otimes R$ for $Q \in \mathcal{P}_m$ and $R \in \mathcal{P}_{n-m}$. Equivalently,

$$G(P \otimes I^{n-m}) = (Q \otimes R)G.$$

It follows that

$$\begin{aligned} H(P|x\rangle) \otimes \left|\psi'\right\rangle &= G(P|x\rangle \otimes |\psi\rangle) \\ &= (Q \otimes R)G(|x\rangle \otimes |\psi\rangle) \\ &= (Q \otimes R)\left(H(|x\rangle) \otimes \left|\psi'\right\rangle\right) \\ &= QH(|x\rangle) \otimes R\left|\psi'\right\rangle, \end{aligned}$$

so up to phase, $H(P|x\rangle) = QH(|x\rangle)$ for all $|x\rangle$, and $|\psi'\rangle = R|\psi'\rangle$. The first equation implies $HP = QH$, or $HPH^\dagger = Q$. Since P was arbitrary, the conjugation of a Pauli string by H is always another Pauli string, so H is a stabilizer gate.

In the special case that P (and therefore $P \otimes I^{n-m}$) is a Pauli basis element, then $Q \otimes R$ is represented in row of the binary tableau of G. We keep the bits representing Q in the tableau for H, since $HPH^\dagger = Q$, which is why $\mathcal{M}(H)$ is a submatrix of $\mathcal{M}(G)$. Clearly the phase of Q is the same as the phase of $Q \otimes R$ if and only if R is positive. In the special case $R = \pm I^{n-m}$, it is easy to see that R is positive, otherwise

$$\left|\psi'\right\rangle = R\left|\psi'\right\rangle = -I^{n-m}\left|\psi'\right\rangle = -\left|\psi'\right\rangle$$

is a contradiction. Hence, the phase for the corresponding row of $\mathcal{T}(H)$ is inherited from $\mathcal{T}(G)$. □

As a direct consequence of these rules, our invariants take on a distinctly algebraic flavor. Let us consider, for the sake of illustration, invariants that depend only on the matrix part of the tableau and ignore the phase bits. Then an invariant is equivalent to a set of matrices closed under the four rules above. In particular, the matrices to form a group under multiplication as a consequence of the composition rule (and the fact that every gate has finite order).

On the other hand, not every group of matrices will correspond to an invariant. For instance, due to the swap rule, the group of matrices must also be closed under arbitrary reordering the rows and columns. This eliminates, e.g., the group of upper triangular matrices. Similarly, the ancilla rule excludes the special orthogonal group. In the end, we are left with just two kinds of matrix groups which lead to invariants:

Subring Invariants Matrices with elements restricted to a particular subring of R (analogous to the real matrices, integer matrices, etc.)

Permutation Invariants Permutation matrices, except where each 1 entry can be any one of a subset of invertible elements, and each 0 entry comes from a collection of non-invertible elements.

Now we are ready to present formal definitions for these invariants, and show that they really are preserved by the circuit-building rules.

6.2 Subring invariants

The first kind of invariant restricts the entries of the tableau to a subring of R. That is, given a subring $\mathsf{S} \subseteq \mathsf{R}$, a gate satisfies the invariant $\mathcal{I}(\mathsf{S})$ if and only if all entries of the tableau are in S. There are twelve classes, all near the top of the lattice, of the form

$$\mathcal{C} = \{\text{All gates satisfying } \mathcal{I}(\mathsf{S})\},$$

corresponding to all 12 subrings of R listed below.

- The entire ring, R, is technically a subring of itself, and $\mathcal{I}(\mathsf{R})$ is the trivial invariant satisfied by all stabilizer gates. Notice that not *every* matrix over R gives a valid tableau because it must still be unitary.

- There are four maximal proper subrings of R:

$$\begin{aligned}
\mathsf{R}_X &= \{\left(\begin{smallmatrix} a & 0 \\ c & d \end{smallmatrix}\right) : \{a, c, d\} \in \{0, 1\}\}, \\
\mathsf{R}_Y &= \{\left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right) : \{a, b, c, d\} \in \{0, 1\}, a + b + c + d = 0\}, \\
\mathsf{R}_Z &= \{\left(\begin{smallmatrix} a & b \\ 0 & d \end{smallmatrix}\right) : \{a, b, d\} \in \{0, 1\}\}, \\
\mathsf{R}_E &= \{\left(\begin{smallmatrix} a & b \\ b & a+b \end{smallmatrix}\right) : \{a, b\} \in \{0, 1\}\}.
\end{aligned}$$

  Our formal definition for Z-preserving gates is the invariant $\mathcal{I}(\mathsf{R}_Z)$. The fact that the lower left entry is 0 implies that the gate maps Pauli strings of I and Z to strings of I and Z. Hence, Z-basis strings are mapped to Z-basis strings. Similarly, the X-preserving and Y-preserving invariants are $\mathcal{I}(\mathsf{R}_X)$ and $\mathcal{I}(\mathsf{R}_Y)$ respectively. The egalitarian invariant, $\mathcal{I}(\mathsf{R}_E)$, comes from the subring R_E.

- The intersection of two subrings is itself a subring, giving us exactly four more subrings ($\mathsf{R}_X \cap \mathsf{R}_Y$, $\mathsf{R}_X \cap \mathsf{R}_Z$, $\mathsf{R}_Y \cap \mathsf{R}_Z$, and $\mathsf{R}_X \cap \mathsf{R}_Y \cap \mathsf{R}_Z$) since the intersection of R_E with any of the others is

  $$\mathsf{R}_X \cap \mathsf{R}_Y \cap \mathsf{R}_Z = \{\left(\begin{smallmatrix} 0 & 0 \\ 0 & 0 \end{smallmatrix}\right), \left(\begin{smallmatrix} 1 & 0 \\ 0 & 1 \end{smallmatrix}\right)\},$$

  the trivial ring.

- Three more subrings are obtained by taking only self-conjugate elements of R_X, R_Y, and R_Z respectively. An element $\left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)$ is *self-conjugate* if

  $$\left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right) = \left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)^*,$$

  or equivalently, $a = d$. These invariants correspond to the X-orthogonal (i.e., $\langle \mathrm{T}_4, \mathcal{P}, \mathrm{R}_X \rangle$), Y-orthogonal (i.e., $\langle \mathrm{T}_4, \mathcal{P}, \mathrm{R}_Y \rangle$), and Z-orthogonal (i.e., $\langle \mathrm{T}_4, \mathcal{P}, \mathrm{R}_Z \rangle$) classes respectively.

Theorem 4. *For any subring $\mathsf{S} \subseteq \mathsf{R}$, the property $\mathcal{I}(\mathsf{S})$ is an invariant. That is, the set of matrices over S respect the circuit building operations.*

Proof. Every subring contains $\left(\begin{smallmatrix} 0 & 0 \\ 0 & 0 \end{smallmatrix}\right)$ and $\left(\begin{smallmatrix} 1 & 0 \\ 0 & 1 \end{smallmatrix}\right)$ by definition, and therefore the tableau (phase bits omitted) of the SWAP gate,

$$\left(\begin{array}{cc|cc} 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ \hline 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{array}\right)$$

satisfies $\mathcal{I}(\mathsf{S})$.

Matrix multiplication is a polynomial in the entries of the two matrices, so composition cannot produce entries outside the subring. Similarly, combining tableau with tensor products or reducing tableau to submatrices via ancillas does not introduce any new ring elements; those operations only use elements already present in the tableau. We conclude that $\mathcal{I}(\mathsf{S})$ is an invariant for any subring S. □

6.3 Permutation invariants

The *permutation invariants* get their name from the matrix part of their tableaux, which is required to have the structure of a permutation matrix. That is, every row (or column) has exactly one element which is invertible, and the others are non-invertible. Permutation invariants are also sensitive to phase bits. It is natural to associate the unique invertible element in a row with the phase bits for that row, giving the tableau of a single-qubit gate. A permutation invariant $\mathcal{P}(G, S)$ is defined by the set of single-qubit gates G which can be obtained in this way, and the set of non-invertible elements S used to fill the rest of the tableau. In other words, a tableau satisfies $\mathcal{P}(G, S)$ if all entries are from S except exactly one entry per row which, when combined with the phase bits for the row, is the tableau of some gate in G.

Note that not all pairs of sets (G, S) produce an invariant. For instance, circuit-building operations will fail to preserve $\mathcal{P}(G, S)$ if G is not a group. The exact relationship between G and S required to produce an invariant is difficult to write down. Roughly speaking, products of elements in S should be zero, products of elements in G should remain in G, and products between S and $\mathcal{M}(G)$ should be manageable in some sense. Theorem 5 gives a list of $\mathcal{P}(G, S)$ invariants, which will turn out to be exhaustive by Theorem 19, the culminating theorem of this paper.

Theorem 5. *We prove that the following permutation invariants are indeed invariant under the circuit-building operations.*

1. *If G is a group of single-qubit gates then*

$$\mathcal{P}(G, \{\left(\begin{smallmatrix}0&0\\0&0\end{smallmatrix}\right)\})$$

is an invariant for $\langle G\rangle$. All thirty degenerate classes are characterized by invariants of this form.

2. *If $\langle X\rangle \subseteq G \subseteq \langle \mathcal{P}, \mathrm{R}_X\rangle$ is a group of single-qubit gates then*

$$\mathcal{P}(G, \{\left(\begin{smallmatrix}0&0\\0&0\end{smallmatrix}\right), \left(\begin{smallmatrix}0&0\\1&0\end{smallmatrix}\right)\})$$

is an invariant for $\langle C(X,X), G\rangle$. These invariants characterize the five X-degenerate classes.

3. *If $\langle Y\rangle \subseteq G \subseteq \langle \mathcal{P}, \mathrm{R}_Y\rangle$ is a group of single-qubit gates then*

$$\mathcal{P}(G, \{\left(\begin{smallmatrix}0&0\\0&0\end{smallmatrix}\right), \left(\begin{smallmatrix}1&1\\1&1\end{smallmatrix}\right)\})$$

is an invariant for $\langle C(Y,Y), G\rangle$. These invariants characterize the five Y-degenerate classes.

4. *If $\langle Z\rangle \subseteq G \subseteq \langle \mathcal{P}, \mathrm{R}_Z\rangle$ is a group of single-qubit gates then*

$$\mathcal{P}(G, \{\left(\begin{smallmatrix}0&0\\0&0\end{smallmatrix}\right), \left(\begin{smallmatrix}0&1\\0&0\end{smallmatrix}\right)\})$$

is an invariant for $\langle C(Z,Z), G\rangle$. These invariants characterize the five Z-degenerate classes.

Proof. Let $\mathcal{P}(G,S)$ be one of the invariants above. Let $M = \{\mathcal{M}(g) : g \in G\}$ be the set of matrices from tableaux in G. In all cases, S contains $\left(\begin{smallmatrix}0&0\\0&0\end{smallmatrix}\right)$, and G contains the single-qubit identity operation,

$$\left(\begin{array}{cc|c} 1 & 0 & 0 \\ 0 & 1 & 0 \end{array}\right),$$

so SWAP satisfies the invariant. And clearly the direct sum of two tableaux in $\mathcal{P}(G,S)$ is still in $\mathcal{P}(G,S)$ for any G and S.

Now consider the composition of two gates. Each entry in the tableau is a dot product of some row from one tableau with some column from the other. Hence, the entry is a sum of $S\times S$, $S\times M$, $M\times S$, or $M\times M$ products. Observe that the $S\times S$ products are all zero (for the particular sets S above), so we may ignore those products. Recall that the row and column each contain exactly one entry in M, so depending on whether those entries align, we get either $SM + MS \subseteq S$ or $M^2 = M$. Furthermore, for any row in one tableau there is exactly one column in the other such that the invertible entries line up. Therefore, exactly one entry in any row (or column) of the composition is in M and the rest are in S. Clearly the matrix part of the tableau has the correct form for the invariant.

We must also consider phase bits under composition. Recall that the phase bits associate with the invertible entries of the matrix to produce single-qubit gates. When we multiply two tableau, these single-qubit gates multiply to produce elements in G (since G is a group), as you would expect. If the non-invertible elements are all zero, then this is the only factor in determining phase bits, so the invariant is preserved by composition.

Now consider the phase bits in the case where S contains nonzero elements, for instance,

$$S = \{\left(\begin{smallmatrix}0&0\\0&0\end{smallmatrix}\right), \left(\begin{smallmatrix}0&1\\0&0\end{smallmatrix}\right)\}.$$

Notice that in this case, both matrices in S have zeros in the bottom row, and the invertible matrices are of the form $\left(\begin{smallmatrix}1&a\\0&1\end{smallmatrix}\right)$. Hence, every even-indexed row of the tableau (as a binary matrix) is all zeros except for one entry. Using the method of tableau composition in Section 4, one can easily show that for these even-indexed rows, the phase bits are exactly what one would get by composing the invertible elements as gates in G. For the other half of the rows, the non-invertible elements may flip the phase bits. But we assume G contains the Pauli element (in this case Z) which flips that sign, so the invertible elements and associated phase bits are still in G, therefore the invariant is preserved. The X- and Y-degenerate cases are similar.

Last, we show that $\mathcal{P}(G, S)$ is preserved under ancilla operations. Recall that when we use ancillas, we remove the rows and columns corresponding to those bits. Clearly the elements of the submatrix are still in M and S. There is a risk that the invertible element for some row could be in one of the removed columns, but if the submatrix is missing an invertible element in some row then the submatrix is not unitary and the ancilla rule must have been misapplied. Hence, only elements in S are removed in the non-ancilla part of the tableau, and each row still contains exactly one entry in M.

We appeal to Theorem 3 for the phase bits. The theorem says that removing the ancillas can only change the sign for a row if there is a nonzero entry in the non-ancilla bits of the row that are removed. For example, if $S = \{\left(\begin{smallmatrix}0&0\\0&0\end{smallmatrix}\right), \left(\begin{smallmatrix}0&1\\0&0\end{smallmatrix}\right)\}$ then only the top phase bit can change. But changing the top phase bit is the same as applying a Z, and for this case Z is assumed to be in G, so the combination of the element in M and the phase bits is still in G. Therefore the Z-degenerate $\mathcal{P}(G, S)$ are invariants, and the X-degenerate and Y-degenerate invariants follow by symmetry. □

7 Equivalence of Generator and Invariant Definitions

We have now defined each class by a set of generators, and by an invariant, but have not yet shown that these definitions coincide. Below are a collection of lemmas which prove this for all classes in our lattice. Note that one direction is always trivial: it is easy to check that the generators defining a class satisfy a particular invariant, and therefore everything they generate (i.e., the class) must satisfy the invariant. We encourage the reader to check these invariants against, say, the tableaux in Table 2.

For the other inclusion (i.e., every gate satisfying the invariant can be generated by the given generators), we start with an arbitrary gate g satisfying the invariant, and apply gates in the class to g to simplify its tableau step-by-step until it is the tableau of the identity operation. It follows that $AgB = I$ for circuits A and B in the class, which proves $g = A^{-1}B^{-1}$ is in the class. In many cases, the circuit derived this way is a *canonical form* for the gate, and can be used to count the number of gates on n qubits in a class.

Let us start with the degenerate classes.

Lemma 6. *Let G be a group of single-qubit gates, and let g be a gate satisfying the permutation invariant $\mathcal{P}(G, \{\left(\begin{smallmatrix}0&0\\0&0\end{smallmatrix}\right)\})$. Then there is a circuit for g consisting of a permutation of the inputs followed by layer of single-qubit gates in G.*

Proof. Consider the tableau for g. Each row or column has exactly one invertible element, so we can read off a permutation π from the positions of those elements. Apply SWAP gates to g to remove this permutation, and put the invertible elements on the diagonal. When we pair a diagonal element with the phase bits for that row, we get a single-qubit gate g_i in G. Applying the inverse of this gate to qubit i will zero the phase bits for that row, and make $\left(\begin{smallmatrix}1&0\\0&1\end{smallmatrix}\right)$ the diagonal entry. Once we do this for each row, we have the identity tableau, therefore g is in $\langle G\rangle$ and has a circuit of the desired form. □

Next, we consider the Z-degenerate classes and, by symmetry, the X-degenerate and Y-degenerate classes.

Lemma 7. *Let $\langle Z\rangle \subseteq G \subseteq \langle \mathcal{P}, \mathrm{R}_Z\rangle$ be a subgroup of Z-preserving single-qubit gates. Let g be any gate satisfying the permutation invariant $\mathcal{P}(G, \{\left(\begin{smallmatrix}0&0\\0&0\end{smallmatrix}\right), \left(\begin{smallmatrix}0&1\\0&0\end{smallmatrix}\right)\})$. Then there is a circuit for g consisting of a layer of single-qubit gates (from G), a layer of $C(Z,Z)$ gates, and a permutation.*

Proof. Consider the tableau of g. We can read off a permutation π, and a single-qubit gate for each input. Assume we have removed those gates (i.e., we now consider the tableau of $g\pi^{-1}g_1^{-1}\cdots g_n^{-1}$), so the tableau has $\left(\begin{smallmatrix}1&0\\0&1\end{smallmatrix}\right)$ on the diagonal, all other entries are either $\left(\begin{smallmatrix}0&1\\0&0\end{smallmatrix}\right)$ or zero, and the phase bits are zero.

The non-zero, off-diagonal entries in the matrix indicate the positions of $C(Z,Z)$ gates. Specifically, if the entry in row i and column j is nonzero then there is a $C(Z,Z)$ on qubits i and j. Note that because the matrix part of the tableau is unitary, the symmetric entry in row j and column i must also be non-zero. The remainder of the circuit consists of the set of $C(Z,Z)$ gates indicated by the non-zero, off-diagonal entries. Notice that $C(Z,Z)$ gates always commute, so their ordering does not matter.. □

Now let us consider four Z-preserving classes which, when we consider symmetry (i.e., the X-preserving and Y-preserving equivalents) cover all but two of the remaining classes.

Lemma 8. *Each of the classes $\langle \mathrm{T}_4, \mathcal{P}\rangle$, $\langle \mathrm{T}_4, \mathcal{P}, \mathrm{R}_Z\rangle$, $\langle C(Z,X), \mathcal{P}\rangle$, and $\langle C(Z,X), \mathcal{P}, \mathrm{R}_Z\rangle$ is the set of all gates corresponding to a subring invariant, where the subrings are*

$$\begin{aligned}
\mathsf{S}_1 &= \{\left(\begin{smallmatrix}0&0\\0&0\end{smallmatrix}\right), \left(\begin{smallmatrix}1&0\\0&1\end{smallmatrix}\right)\},\\
\mathsf{S}_2 &= \{\left(\begin{smallmatrix}0&0\\0&0\end{smallmatrix}\right), \left(\begin{smallmatrix}0&1\\0&0\end{smallmatrix}\right), \left(\begin{smallmatrix}1&0\\0&1\end{smallmatrix}\right), \left(\begin{smallmatrix}1&1\\0&1\end{smallmatrix}\right)\},\\
\mathsf{S}_3 &= \{\left(\begin{smallmatrix}0&0\\0&0\end{smallmatrix}\right), \left(\begin{smallmatrix}1&0\\0&0\end{smallmatrix}\right), \left(\begin{smallmatrix}0&0\\0&1\end{smallmatrix}\right), \left(\begin{smallmatrix}1&0\\0&1\end{smallmatrix}\right)\},\\
\mathsf{S}_4 &= \{\left(\begin{smallmatrix}0&0\\0&0\end{smallmatrix}\right), \left(\begin{smallmatrix}0&0\\0&1\end{smallmatrix}\right), \left(\begin{smallmatrix}0&1\\0&0\end{smallmatrix}\right), \left(\begin{smallmatrix}0&1\\0&1\end{smallmatrix}\right), \left(\begin{smallmatrix}1&0\\0&0\end{smallmatrix}\right), \left(\begin{smallmatrix}1&0\\0&1\end{smallmatrix}\right), \left(\begin{smallmatrix}1&1\\0&0\end{smallmatrix}\right), \left(\begin{smallmatrix}1&1\\0&1\end{smallmatrix}\right)\}.
\end{aligned}$$

Proof. In all four classes, elements of the tableau are of the form $\left(\begin{smallmatrix}a&b\\0&d\end{smallmatrix}\right)$. Suppose $\left(\begin{smallmatrix}?&?\\0&d_1\end{smallmatrix}\right)$ is the ith entry of some row, and $\left(\begin{smallmatrix}?&?\\0&d_2\end{smallmatrix}\right)$ is the jth entry in the same row, where entries labeled by "?" are unconstrained. If we apply a CNOT gate from qubit i to j, these entries will be of the form $\left(\begin{smallmatrix}?&?\\0&d_1\end{smallmatrix}\right)$ and $\left(\begin{smallmatrix}?&?\\0&d_1+d_2\end{smallmatrix}\right)$ respectively. That is, the bottom right bits change as though we applied the CNOT gate to those bits. Since a T_4 gate can be built from CNOT gates, it will (similarly) affect the bottom right bits as though we are applying a T_4.

Our strategy is to use either CNOT or T_4 gates (depending on the class) to perform Gaussian elimination on the bottom right entries of the matrix elements. If we have access to CNOT gates then we literally apply Gaussian elimination, using CNOT to add one column to another, and using SWAP to exchange columns.

If we only have T_4 gates then we are in subring $\mathsf{S}_1 \subseteq \mathsf{S}_2$ or S_2, so $\left(\begin{smallmatrix}1 & 0\\ 0 & 1\end{smallmatrix}\right)$ and $\left(\begin{smallmatrix}1 & 1\\ 0 & 1\end{smallmatrix}\right)$ are the only elements with a 1 in the bottom right position, and also the only invertible elements. It follows that the number of bottom right bits set to 1 in a row is the same as the number of invertible elements, which must be odd because the matrix is unitary. To reduce the number, we apply a T_4 to three 1 bits and a 0 bit (note: we may add a zero bit by adding an ancilla, if necessary), which changes the 0 to a 1 and the 1's to 0's, reducing the number of 1's (or invertible elements) in the row by two. When there is a single 1 left in the row, unitarity conditions imply that it is also the only 1 left in that column, so we may ignore that row and column for the moment and continue to eliminate the rest of the matrix.

Now suppose we have row reduced the matrix, using either CNOT or T_4, so that the bottom right entry of every element is 0, except along the diagonal where that bit 1. At this point, the diagonal element is the only element in a row that can possibly be invertible, therefore the diagonal elements are of the form $\left(\begin{smallmatrix}1 & b\\ 0 & 1\end{smallmatrix}\right)$. Similarly, unitarity conditions imply that the off-diagonal elements are of the form $\left(\begin{smallmatrix}0 & b\\ 0 & 0\end{smallmatrix}\right)$. In other words, the remaining tableau is Z-degenerate, since there is only one invertible element per row or column, and the off-diagonal elements are in $I = \{\left(\begin{smallmatrix}0 & 0\\ 0 & 0\end{smallmatrix}\right), \left(\begin{smallmatrix}0 & 1\\ 0 & 0\end{smallmatrix}\right)\}$. We can use either Lemma 6 or Lemma 7 to find a circuit from the remainder, which is in either $\langle \mathcal{P} \rangle$ or $\langle C(Z,Z), \mathcal{P}, \mathrm{R}_Z \rangle$, depending on the class. □

There are only two classes remaining, ALL and $\langle \mathrm{T}_4, \mathcal{P}, \Gamma \rangle$, which we handle specially. For the first, we appeal to Aaronson and Gottesman [2] who give an explicit decomposition for any stabilizer gate into layers of CNOT, Hadamard (θ_{X+Z} in our notation), and phase (R_Z) gates.

Lemma 9. *Any egalitarian gate g can be constructed from T_4, $\mathcal{P}$ and Γ gates.*

Proof. Egalitarian gates satisfy the invariant that all elements are in the subring

$$\{\left(\begin{smallmatrix}0 & 0\\ 0 & 0\end{smallmatrix}\right), \left(\begin{smallmatrix}1 & 0\\ 0 & 1\end{smallmatrix}\right), \left(\begin{smallmatrix}1 & 1\\ 1 & 0\end{smallmatrix}\right), \left(\begin{smallmatrix}0 & 1\\ 1 & 1\end{smallmatrix}\right).\}$$

In fact, this subring is isomorphic to $\mathbb{F}_4$, so it is a field. The unitarity of the matrix in our sense translates to unitarity as a matrix over $\mathbb{F}_4$.

Like the other T_4 classes, we use Gaussian elimination on the tableau of g. Consider a row of the tableau. If the entry in some column is not the identity, then apply Γ or Γ^{-1} to the corresponding qubit to make it the identity. By unitarity, there are an odd number of identity elements in the row. We may remove pairs of identity elements with a T_4, similar to Lemma 8, until there is only one left and the rest of the row is zero. Unitarity implies the column below the identity element is also zero, and we proceed to eliminate the rest of the tableau. Once the matrix part of the tableau is the identity, we apply Pauli matrices to zero out the phase bits.

We conclude that all egalitarian gates are in $\langle \mathrm{T}_4, \mathcal{P}, \Gamma \rangle$. □

8 Circuit Identities

In this section, we give necessary tools to prove that a set of gates generates, in some sense, "all that one could hope for." Formally, we wish to prove that the gate set generates a particular class in the classification lattice when it is contained in that class but fails to satisfy the invariants of all classes below it. To this end, we give several useful circuit identities that will be used extensively in Section 9. For instance, one can show that any circuit on two qubits can be reduced to an equivalent circuit containing at most one generalized CNOT gate (see Appendix D). The following

lemma gives only the aspect of that theorem that is necessary to the classification, that is, the ability to extract single-qubit stabilizers from the composition of generalized CNOT gates.

Lemma 10. *Let $P, Q, R \in \mathcal{P}$, and let $\Gamma P \Gamma^\dagger = Q$ and $\Gamma Q \Gamma^\dagger = R$. Then*

- *$C(P,Q)$ and $C(P,R)$ generate R_P.*
- *$C(P,P)$ and $C(P,R)$ generate R_P.*
- *$C(P,P)$ and $C(Q,R)$ generate Γ.*
- *$C(P,P)$ and $C(Q,Q)$ generate θ_{P+Q}.*

Proof. The first two inclusions come from the following identity, which holds whenever $\Gamma Q \Gamma^\dagger = R$ (i.e. regardless of P):

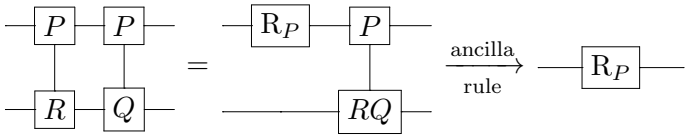

Similarly, for the third identity, we get

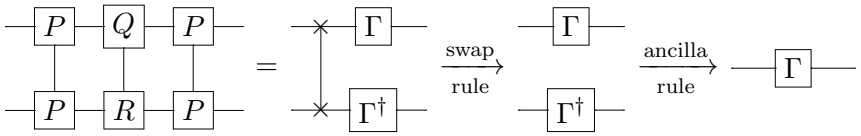

and for the final identity

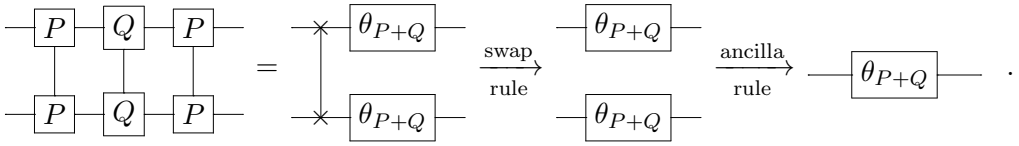

□

It might seem strange to reduce non-degenerate gates into less powerful single-qubit gates, but we will eventually see that single-qubit generators are the most crucial. Once we have shown that a particular set of gates generates all single-qubit operations, then that set of gates will generate the class of *all* stabilizer operations provided it contains any non-degenerate gate. All non-degenerate gates generate at least one Pauli, often the entire Pauli group, which is why some single-qubit classes do not appear as the single-qubit subgroup of a non-degenerate class. For instance, consider the CNOT gate where the first qubit controls the second qubit. If we let the first input be $|1\rangle$, then a Pauli X operation is always applied to the second qubit. Similarly, if we let the input to the second qubit be $|-\rangle$, then a Pauli Z operation is always applied to the first qubit. Under the ancilla rule, we now have Pauli X and Z operations, so we can generate Y and the entire Pauli group. Clearly, the same is true for any heterogeneous CNOT gate. However, surprisingly, the following lemma shows that even the T_4 gate suffices to generate the entire Pauli group.

Lemma 11. T_4 *generates the Pauli group.*

Proof. Consider the following two circuits:

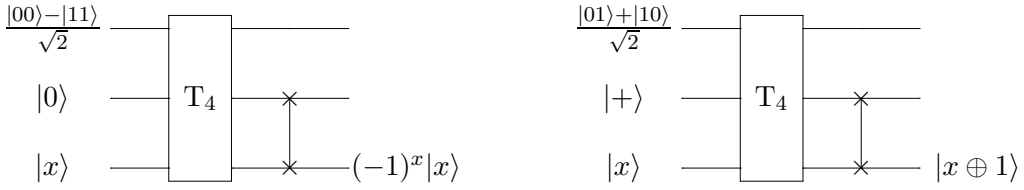

Under the ancilla rule, the first generates a Pauli Z operation while the second generates a Pauli X, from which we can clearly generate the Pauli group. □

There is another way to view the identity of Lemma 11 which will be useful later. Since T_4 is an affine gate over the computational basis states, $\mathrm{T}_4 = C_1 C_2 \dots C_n$ where each C_i is a CNOT gate. Furthermore, T_4 and CNOT are their own inverses, so $\mathrm{T}_4 = C_n C_{n-1} \dots C_1$. Finally, because T_4 is symmetric when represented as a 4×4 matrix over $\mathbb{F}_2$, $\mathrm{T}_4 = C_1^T C_2^T \dots C_n^T$. Notice that C_i^T just represents the CNOT gate C_i where the control and target qubits are swapped. Therefore, leveraging the well-known equivalence $\theta_{X+Z}^{\otimes 2} \circ \mathrm{CNOT} \circ \theta_{X+Z}^{\otimes 2} = \mathrm{SWAP} \circ \mathrm{CNOT} \circ \mathrm{SWAP}$ we arrive at the following consequence:[6]

$$\theta_{X+Z}^{\otimes 4}\, \mathrm{T}_4\, \theta_{X+Z}^{\otimes 4} = \mathrm{T}_4\,.$$

Similarly, by straightforward calculation, we get

$$\mathrm{R}_Z^{\otimes 4}\, \mathrm{T}_4\, \mathrm{R}_Z^{\otimes 4} = \mathrm{T}_4\,.$$

It is now easy to extend old circuit identities into new ones. For instance, conjugating the first circuit in the proof of Lemma 11 by $\theta_{X+Z}^{\otimes 4}$ (which does not change the circuit because of the above observations) and pushing the θ_{X+Z} gates into the inputs, yields the second circuit. This technique is in fact very general and is used in the proof of the lemma below.

Lemma 12. T_4 *and* $C(P,P)$ *generate* R_P.

Proof. Figure 6 shows how to generate R_Z with $C(Z,Z)$. Using the argument by conjugation above, T_4 and $C(P,P)$ generate R_P. □

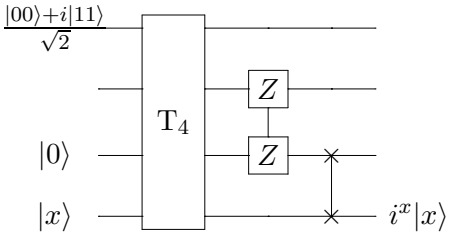

Figure 6: Generating R_Z with T_4 and $C(Z,Z)$.

The following lemmas make precise our working assumption that single-qubit gates can significantly bolster the power of non-degenerate gate sets.

Lemma 13. *Suppose we have any* $C(P,Q)$ *gate with any single-qubit gate* G *that does not preserve the* P*-basis and any single-qubit gate* H *that does not preserve the* Q*-basis. Then* $\langle C(P,Q), G, H\rangle = \mathsf{ALL}$.

[6] Recall that θ_{X+Z} is commonly known as the Hadamard gate.

Proof. We will prove that the class $\langle C(P,Q), G, H\rangle$ contains all single-qubit gates. Then, to prove that the class generates all stabilizer operations, it is be sufficient to show that it contains a CNOT gate. However, since all generalized CNOT gates are conjugates of each other, this is immediate.

First suppose $P = Q$. Since G does not preserve P-basis, we can use G to create a $C(R,R)$ gate where $R \neq P$. By Lemma 10, we can generate a θ_{P+R} gate. Conjugating $C(P,P)$ by θ_{P+R} on the second qubit yields a $C(P,R)$ gate. Once again leveraging Lemma 10, $C(P,R)$ and $C(P,P)$ generate an R_P gate. Referring to the single-qubit lattice (see Figure 4), we see that the class $\langle \mathcal{P}, \theta_{P+R}, \mathrm{R}_P\rangle$ contains all single-qubit gates.

Now suppose that $P \neq Q$. Once again, since G does not preserve P-basis, we can use G to create a $C(R,Q)$ gate. If $R = Q$, then by the logic above, we can use H to generate all single qubit gates, so suppose $R \neq Q$. By Lemma 10, we can use $C(P,Q)$ and $C(R,Q)$ to generate an R_Q gate. Conjugating both $C(P,Q)$ and $C(R,Q)$ by H appropriately, gives a $C(P,S)$ and $C(R,S)$ for some $S \neq Q$, which we can once again generate an R_S gate. Referring to the single-qubit lattice, we see that the class $\langle \mathcal{P}, \mathrm{R}_S, \mathrm{R}_Q\rangle$ contains all single-qubit gates. □

Lemma 14. T_4 *with the class of all single-qubit gates generates* ALL.

Proof. It is well known that CNOT, θ_{X+Z}, and R_Z generate all stabilizer circuits. Therefore, it will be sufficient to show that T_4 plus all single-qubit gates generate CNOT. Under the ancilla rule, it is clear by Figure 7 that T_4 and R_Z suffice to generate $C(Z,Z)$. Conjugating one qubit of $C(Z,Z)$ by θ_{X+Z} yields a $C(Z,X) = \text{CNOT}$ gate, completing the proof. □

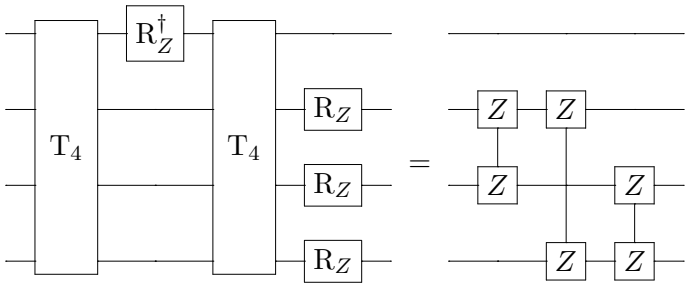

Figure 7: Generating $C(Z,Z)$ with T_4 and R_Z.

9 Universal Construction

Suppose G is an n-qubit stabilizer gate. It turns out there is a single circuit $\mathfrak{C}(G)$, the *universal construction*, which can help us extract useful generators (e.g., single-qubit gates, generalized CNOTs, etc.) from G. Specifically, the circuit $\mathfrak{C}(G)$ (shown in Figure 8) applies G to qubits 2 through $n+1$, swaps qubits 1 and 2, then applies G^{-1} to qubits 2 through $n+1$.

Note that after we apply G, all of the qubits but one go directly into G^{-1}, which should intuitively cancel out "most" of the effect G has on those qubits, isolating the effect of G on the swapped qubit. The following theorem makes this intuition more precise.

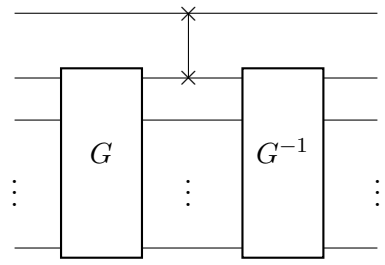

Figure 8: Universal Construction $\mathfrak{C}(G)$

Theorem 15. *Let G be an n-qubit stabilizer gate. Then*

$$\mathcal{M}(\mathfrak{C}(G)) = I_{n+1} + \begin{pmatrix} 1 \\ v \end{pmatrix} \begin{pmatrix} 1 & v^* \end{pmatrix}$$

where $v \in \mathsf{R}^{n\times 1}$ is the first column of $\mathcal{M}(G)$.

Proof. Let $\mathcal{M}(G) = \left(\begin{smallmatrix} A & B \\ C & D \end{smallmatrix}\right)$ with $A \in \mathsf{R}$, $B \in \mathsf{R}^{1\times(n-1)}$, $C \in \mathsf{R}^{(n-1)\times 1}$, and $D \in \mathsf{R}^{(n-1)\times(n-1)}$. Thus,

$$\mathcal{M}(G)\mathcal{M}(G^\dagger) = \begin{pmatrix} 1 & 0 \\ 0 & I_{n-1} \end{pmatrix} = \begin{pmatrix} A & B \\ C & D \end{pmatrix} \begin{pmatrix} A^* & C^* \\ B^* & D^* \end{pmatrix} = \begin{pmatrix} AA^* + BB^* & AC^* + BD^* \\ CA^* + DB^* & CC^* + DD^* \end{pmatrix}$$

which implies

$$\begin{aligned} \mathcal{M}(\mathfrak{C}(G)) &= \begin{pmatrix} 1 & 0 & 0 \\ 0 & A & B \\ 0 & C & D \end{pmatrix} \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 \\ 0 & A^* & C^* \\ 0 & B^* & D^* \end{pmatrix} = \begin{pmatrix} 0 & A^* & C^* \\ A & BB^* & BD^* \\ C & DB^* & DD^* \end{pmatrix} \\ &= \begin{pmatrix} 0 & A^* & C^* \\ A & 1 + AA^* & AC^* \\ C & CA^* & I_{n-1} + CC^* \end{pmatrix} = I_{n+1} + \begin{pmatrix} 1 \\ A \\ C \end{pmatrix} \begin{pmatrix} 1 & A^* & C^* \end{pmatrix}. \end{aligned}$$

□

Tableaux of the form $I_n + vv^*$ have relatively simple circuit decompositions in terms of single-qubit generators and generalized CNOT gates, which is formalized in the following theorem.

Theorem 16. *Let $v = \{1, a_2, a_3, \ldots, a_{2k}, b_1, \ldots, b_\ell\} \in \mathsf{R}^{n\times 1}$ where each a_i is invertible and each b_i is singular, and let C be a stabilizer circuit such that $\mathcal{M}(C) = I_n + vv^*$. Then C is equivalent to the circuit consisting of*

- *a T_{2k} gate on the first $2k$ qubits,*
- *conjugated by $\mathcal{G}(a_i, i)$ for all i in $\{2, \ldots, 2k\}$,*
- *conjugated by $\mathrm{CNOT}(b_i, 1, 2k+i)$ for all $i \in \{1, \ldots, \ell\}$, and*
- *a final layer of Pauli gates (*not *conjugated) on every qubit.*

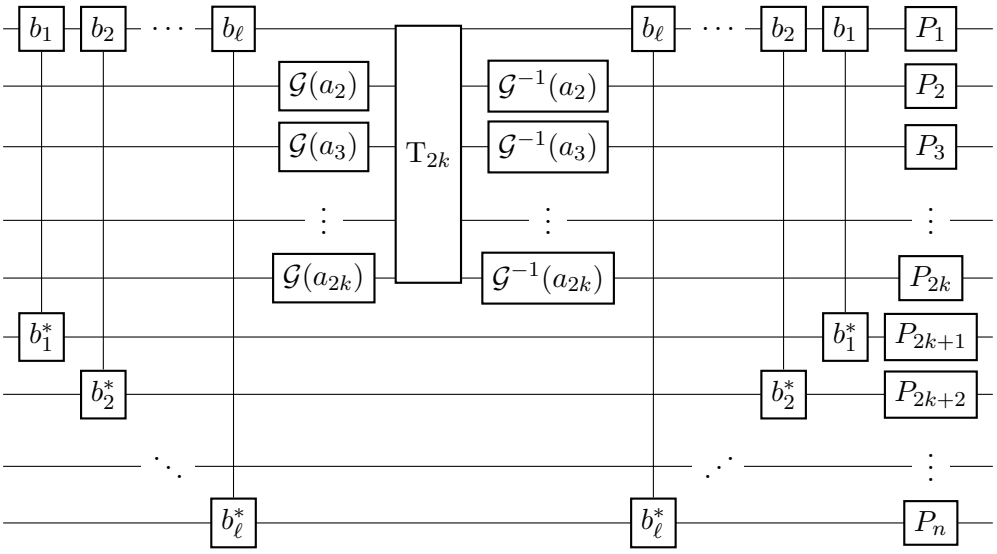

Figure 9: Decomposition of $\mathfrak{C}(G)$.

That is, C is equivalent to the circuit in Figure 9.

Proof. Notice first that because each a_i is invertible, we can conjugate C by $\mathcal{G}(a_i, i)$ for each $i \in \{2, \ldots, 2k\}$, yielding a circuit with the simpler tableau

$$\mathrm{Diag}(1, a_2, \ldots, a_{2k}, 1, \ldots, 1)(I_n + vv^*)\mathrm{Diag}(1, a_2^*, \ldots, a_{2k}^*, 1, \ldots, 1) = I_n + v'v'^*$$

where $v' = \{1, \ldots, 1, b_1, \ldots, b_\ell\}$. Furthermore, conjugating the circuit by the gate $\mathrm{CNOT}(b_i, 1, i + 2k)$ corresponds to the simplification

$$\mathrm{CNOT}(b_i, 1, i + 2k)(I_n + v'v'^*)\,\mathrm{CNOT}(b_i, 1, i + 2k) = I_n + v''v''^*$$

where v'' is equal to v' with the exception that entry $i+2k$ is equal to zero. Repeating this procedure for each $i \in \{1, \ldots, \ell\}$, we arrive at a circuit with a very simple tableau:

$$I_n + \{1, \ldots, 1, 0, \ldots, 0\}\{1, \ldots, 1, 0, \ldots, 0\}^T,$$

which is exactly the tableau of a T_{2k} gate applied to the first $2k$ qubits. Notice, finally, that by reversing the procedure and applying the appropriate Pauli gates to each qubit, we can ensure that the tableau of the decomposition is that of $\mathfrak{C}(G)$. □

Theorem 16 leads to a clean circuit decomposition of $\mathfrak{C}(G)$. All that is left to show is that we can actually generate each of the elementary gates that appears in the decomposition under the ancilla rule. First, we will need the following useful lemma, which will allow us to essentially disregard the Pauli operators in the decomposition of the universal construction when applying the ancilla rule.

Lemma 17. *Let G be a gate on n qubits which is stabilized by some state $|a\rangle$ on the first k qubits and generates H on the remaining qubits. Furthermore, let P be any gate on the first k qubits. Then $P \circ G$ generates $H \otimes H^{-1}$.*

Proof. By supposition we have that $G(|a\rangle \otimes |\psi\rangle) = |a\rangle \otimes H|\psi\rangle$. Therefore $P \circ G(|a\rangle \otimes |\psi\rangle) = P|a\rangle \otimes H|\psi\rangle$. Now apply the inverse of $P \circ G$ to $P|a\rangle \otimes H|\psi\rangle$ with the same first k qubits and $n-k$ new qubits. In the middle, P cancels with P^{-1}, and we can remove those k qubits by using $|a\rangle$ as an ancilla. On the remaining qubits, we have $H \otimes H^{-1}$. □

We are finally ready to prove the main theorem of this section.

Theorem 18. *Let G be a stabilizer gate on n qubits. Furthermore, let*

$$v = (1, a_2, a_3, \ldots, a_{2k}, b_1, \ldots, b_\ell) \in \mathsf{R}^{n \times 1}$$

be a vector where each a_i is invertible and each b_i is singular be some row of $\mathcal{M}(G)$. Then G generates a gate $\mathcal{G}$ such that $\mathcal{M}(\mathcal{G}) = a_i$ for each $i \in \{2, \ldots, 2k\}$, $\mathrm{CNOT}(b_j)$ for all $j \in \{1, \ldots, \ell\}$, and T_{2k} gate.

Proof. From theorems 15 and 16, we know that universal construction $\mathfrak{C}(G)$ can be decomposed as shown in Figure 9. The proof will proceed in the following manner. Starting with the decomposition of $\mathfrak{C}(G)$, we show that it generates some elementary gate. We then use that gate to simplify the original decomposition of $\mathfrak{C}(G)$, eventually generating all such gates in this manner.

First notice that for some input $i \in \{2k+1, \ldots, n\}$, the single-qubit stabilizer state $|b_i\rangle$ serves to remove the effect of the generalized CNOT. By Lemma 17, we can remove the last ℓ qubits with ancillas, at the expense of creating another (inverted) copy of the remaining $2k$ bits. That is, we have $H \otimes H^{-1}$ where H is as follows:

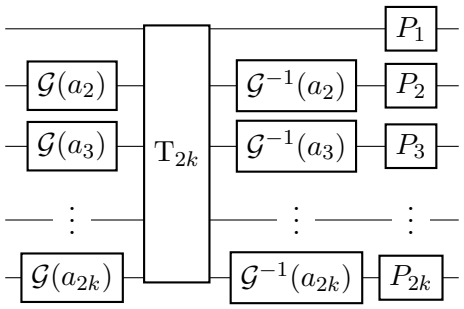

Now let $|\phi\rangle = \frac{|00\rangle+|11\rangle}{\sqrt{2}}$ be the Bell state on two qubits. Notice that we can use $|\phi\rangle$ as an ancilla to remove two bits from $\mathrm{T}_{2\ell}$ (i.e., leaving a $\mathrm{T}_{2\ell-2}$). However, the $\mathrm{T}_{2\ell}$ occurs in H conjugated by single-qubit gates, followed by Pauli operations. If we feed the state $\mathcal{G}(a_i, i)^\dagger \otimes \mathcal{G}(a_j, j)^\dagger |\phi\rangle$ to bits i and j, the single-qubit gates transform it to $|\phi\rangle$, it removes two bits from the $\mathrm{T}_{2\ell}$, then it is transformed to $P_i\mathcal{G}(a_i, i)^\dagger \otimes P_j\mathcal{G}(a_j, j)^\dagger |\phi\rangle$. We can do the exact same thing to H^{-1}, starting with $P_i\mathcal{G}(a_i, i)^\dagger \otimes P_j\mathcal{G}(a_j, j)^\dagger |\phi\rangle$ and going to $\mathcal{G}(a_i, i)^\dagger \otimes \mathcal{G}(a_j, j)^\dagger |\phi\rangle$. Then we swap the two states, and use the ancilla rule to remove them. The net result is that we can remove any two qubits of H, as long as we remove the same pair of qubits from H^{-1}.

Iterate this procedure to until H has been reduced to just the first two qubits, and so has H^{-1}. In particular, the $\mathrm{T}_{2\ell}$ gate in the middle is now a T_2, which we observe is actually a SWAP gate. From H, the remaining circuit is

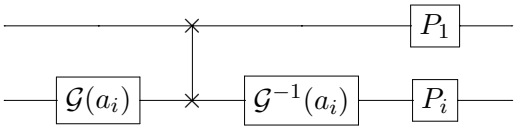

and the inverse survives from H^{-1}. Remove the swaps, and observe that we have ancillas to remove any of the remaining single-qubit gates, so we can isolate each single-qubit gate, in particular $P_i \circ \mathcal{G}^{-1}(a_i)$, $P_1 \circ \mathcal{G}^{-1}(a_i)$ and their inverses. From this point on, everything we do for H can be repeated for H^{-1}, which culminates in removing all the single-qubit gates, at which point we can remove the extra $\mathrm{T}_{2\ell}$ gate easily. Hence, let us ignore H^{-1}.

Now let us repeat the procedure above starting from the $\mathfrak{C}(G)$, but stop short of applying the ancilla rule to qubit $2k + j$. The result is the first circuit depicted below, which is then simplified by an application of the swap rule, and gates $P_i \circ \mathcal{G}^{-1}(a_i)$ and $P_1 \circ \mathcal{G}^{-1}(a_i)$:

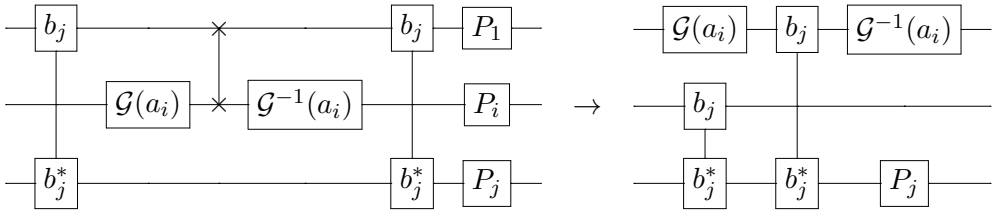

Notice that the topmost qubit is stabilized by $\mathcal{G}^{-1}(a_i)|b_j\rangle$, from which we can see that the ancilla rule immediately generates $\mathrm{CNOT}(b_j)$. Finally, we exploit the identity,

$$T_{2k}(P \otimes I^{\otimes 2k-1})T_{2k} = I \otimes P^{\otimes 2k-1},$$

which holds up to a global phase. We have the following chain of consequences.

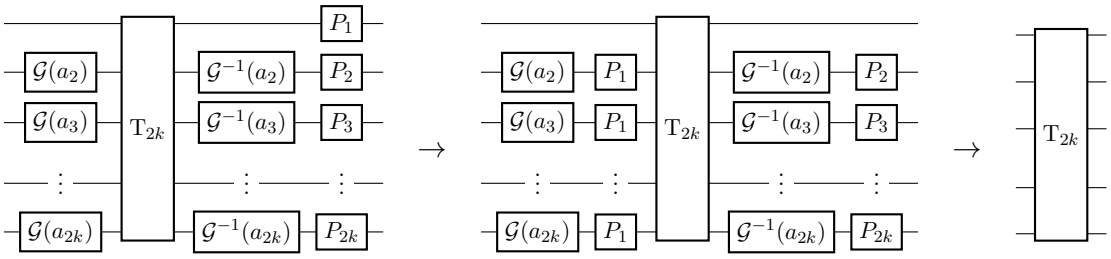

The last implication comes from the fact that we already generated $P_i \circ \mathcal{G}^{-1}(a_i)$ and $P_1 \circ \mathcal{G}^{-1}(a_i)$. □

10 Completing the Classification

The final step in the classification is to demonstrate that the classes we have defined are in fact the only classes that exist.

Theorem 19. *Let $\mathcal{S}$ be some class in the classification, and let G be a collection of gates. Suppose $\langle G \rangle \subseteq \mathcal{S}$, but $\langle G \rangle \not\subseteq \mathcal{S}'$ for all $\mathcal{S}'$ below $\mathcal{S}$ in our classification. That is, for all such $\mathcal{S}'$ there exists a gate $g \in G$ such that $g \notin \mathcal{S}'$. Then $\langle G \rangle = \mathcal{S}$.*

Proof. Let $\mathcal{C}$ be the class generated by G. There is a very general strategy for proving that the given class $\mathcal{C}$ is indeed one already stated in the classification. For each invariant described in Section 6 that $\mathcal{C}$ fails to preserve, the universal construction generates a (simple) gate which also fails to satisfy that invariant by Theorem 18. Composing these gates using identities from Section 8, one can show that they always generate some class in the classification. This would complete the classification.

Nevertheless, we now give a complete sequence of tests to identify the class $\mathcal{C}$. First consider the degeneracy invariant. If $\mathcal{C}$ is degenerate, then by Lemma 6 we can decompose each gate into a circuit of single-qubit gates and swap gates. Each single-qubit gate can be extracted with an appropriate tensor product of single-qubit ancillas on the rest. The question therefore reduces to the simple group-theoretic question about the subgroups of single-qubit gates, which can be solved straightforwardly. Therefore, assume $\mathcal{C}$ is non-degenerate.

Let us separate the remainder of the proof based on the X-, Y-, Z-preserving invariant. Let P, Q, and R be distinct Pauli operations.

Suppose first that $\mathcal{C}$ is X-, Y-, and Z-preserving. Because every generalized CNOT gate violates one such invariant and $\mathcal{C}$ is non-degenerate, some gate in $\mathcal{C}$ must have a tableau with multiple invertible elements in some row. Therefore, from the universal construction we extract a T_4 gate. By Lemma 11, we have then that $\mathcal{C} = \langle \mathrm{T}_4, \mathcal{P} \rangle$.

Suppose now that $\mathcal{C}$ is P- and Q-preserving but not R-preserving. Because $\mathcal{C}$ is non-degenerate, it must generate either a heterogeneous CNOT gate, a homogenous CNOT gate, or a T_4 gate. We wish to show that $\mathcal{C}$ contains a $C(P,Q)$ gate, which would imply $\mathcal{C} = \langle C(P,Q), \mathcal{P} \rangle$ as desired. First notice that no homogeneous CNOT gate can be P- and Q-preserving. Suppose then that from the universal construction, $\mathcal{C}$ generates a T_4 gate but no $C(P,Q)$ gate. Since T_4 is R-preserving, there must be some single-qubit gate from the universal construction that is not R-preserving but is P- and Q-preserving. It is straightforward to check that no such single-qubit gate exists, which implies that $\mathcal{C}$ must contain a $C(P,Q)$ gate.

Suppose now that $\mathcal{C}$ is P-preserving but not Q- and R-preserving. This is the most involved case and will require several more subdivisions. It will be first useful to notice that all P-preserving single-qubit gates are also P-orthogonal. Therefore, if $\mathcal{C}$ violates the P-orthogonality invariant, then the universal construction must produce some non-degenerate gate which violates P-orthogonality. Therefore, $\mathcal{C}$ contains a $C(P,Q)$ gate or a $C(P,R)$ gate. If it contains both via the universal construction, then indeed $\mathcal{C} = \langle C(P,Q), \mathrm{R}_P, \mathcal{P} \rangle$ by Lemma 10. Otherwise, the universal construction produces some single-qubit gate which is P-preserving but not Q- and R-preserving. Since all heterogeneous CNOT gates generate the Pauli group as well, the class of single qubit gates must therefore contain an R_P gate.

Therefore let us now assume that $\mathcal{C}$ is P-orthogonal but not P-degenerate. Since $\mathcal{C}$ is P-orthogonal, it cannot contain a $C(P,Q)$ or $C(P,R)$. However, because $\mathcal{C}$ is P-preserving, it must contain a T_4 gate; otherwise it would be P-degenerate. Once again, since the T_4 gate generates the Pauli group and $\mathcal{C}$ is neither Q- nor R-preserving, the class of single qubit gates must contain an R_P gate, implying that $\mathcal{C} = \langle \mathrm{T}_4, \mathrm{R}_P, \mathcal{P} \rangle$.

Let us then assume that $\mathcal{C}$ is P-degenerate. Since $\mathcal{C}$ is non-degenerate, it must contain a $C(P,P)$ gate. There are five P-degenerate classes, which are determined by their single-qubit subgroup. Indeed, the five P-degenerate classes correspond to the five P-preserving single-qubit classes containing P. Unlike previous cases, such a diversity of classes exists because $C(P,P)$ does not suffice to generate the Pauli group on its own. Once again, the universal construction allows us to extract the $C(P,P)$ gate along with single-qubit gates which suffice to generate every gate in $\mathcal{C}$. It is straightforward to see why the entire single-qubit group must arise from the universal construction. If not, then each gate in $\mathcal{C}$ could be constructed by gates in a smaller class, a contradiction. This completes the classification of all gates that are P-preserving but not Q- and R-preserving.

Assume then that $\mathcal{C}$ is neither P-,Q-, nor R-preserving. If it is egalitarian, then it must contain

a T_4 gate because it is non-degenerate. Therefore it must contain a single-qubit gate that is egalitarian, but not P-,Q-, or R-preserving. The only such single-qubit class which contains the Pauli group is $\langle \Gamma, \mathcal{P} \rangle$. Therefore $\mathcal{C} = \langle \mathrm{T}_4, \Gamma, \mathcal{P} \rangle$.

Finally, let us then assume that $\mathcal{C}$ also violates the egalitarian invariant. That is, $\mathcal{C}$ violates every invariant so should be equal to the class of all stabilizer operations. Suppose the only non-degenerate gate generated by the universal construction is the T_4 gate. In particular, this implies that class of single-qubit gates generated by $\mathcal{C}$ must not be X-, Y-, Z-preserving, nor egalitarian. Therefore the single-qubit class of $\mathcal{C}$ must contain *all* single-qubit gates. Therefore by Lemma 14 we generate all stabilizer circuits.

Similarly, Lemma 10 implies that if the set of generalized CNOT gates generated by the universal construction fails to be P-, Q-, or R-preserving, then $\mathcal{C}$ generates a single-qubit gate which fails to be P-, Q-, or R-preserving, respectively. Combining this fact with Lemma 13 implies that $\mathcal{C} = \mathsf{ALL}$. □

Corollary 20. *Given any set of gates G, there is a subset $S \subseteq G$ of at most three gates such that $\langle S \rangle = \langle G \rangle$.*

Proof. The result follows by a careful accounting of the gates used in the proof of Theorem 19. We give the argument only for the degenerate classes.

Let G be a set of single-qubit gates with $S \subseteq G$ and such that $\langle S \rangle = \langle G \rangle$. Suppose $S = \{g_1, \ldots, g_k\}$ with $k > 3$. We can assume that each generator in the set S is not contained in the subgroup generated by the other elements in S, otherwise we could remove that generator, reducing the size of S. Therefore, there is an ascending chain of subgroups $G_0 \subset G_1 \subset \ldots \subset G_k$, where $G_i = \langle g_1, \ldots, g_i \rangle$. Observing Figure 4, the longest ascending chain has length four, and in particular, all chains of length four end at the class of *all* single-qubit gates. Furthermore, that chain must contain one of the subgroups $\langle \mathcal{P}, \mathrm{R}_X \rangle$, $\langle \mathcal{P}, \mathrm{R}_Y \rangle$, or $\langle \mathcal{P}, \mathrm{R}_Z \rangle$, which we can assume is $\langle \mathcal{P}, \mathrm{R}_Z \rangle$ by symmetry.

Since not all gates in S are Pauli operators, we can assume that g_1 is not a Pauli operator. In particular, $g_1 \in \langle \mathcal{P}, \mathrm{R}_Z \rangle \setminus \langle \mathcal{P} \rangle$. That is, g_1 is one of R_Z, $\mathrm{R}_Z^\dagger$, θ_{X+Y}, or θ_{X-Y}. Therefore, the class G_2 is either $\langle \mathrm{R}_Z \rangle$ or $\langle Z, \theta_{XY} \rangle$. Clearly, $g_4 \notin \langle \mathcal{P}, \mathrm{R}_Z \rangle$ (i.e., it does not preserve the Z-basis), and by a simply case analysis, it is easy to see that $\langle g_1, g_2, g_4 \rangle$ is the entire class of single-qubit operations. □

11 Open Problems

Our classification of stabilizer gates resolves an open problem of Aaronson et al. [3], but leaves their central question, the classification of arbitrary quantum gates, completely open. It is unclear whether there is another piece of the full quantum gate classification that can be peeled off. Other discrete quantum gate sets are known, but none are known to have the rich structure and entanglement of stabilizer gates (aside from conjugated stabilizer gates). So we ask: are there other interesting discrete gate sets, and can they be classified like stabilizer gates?

Another source of open problems is the choice of ancilla rule. As discussed, we permit ancillas initialized to quantum states. We have determined that the classification continues to hold under a stabilizer ancilla model if the following conjecture holds:

Conjecture 21. *For any single-qubit stabilizer gate g, there exists a stabilizer state $|\psi\rangle$ and circuit of* SWAP *gates π such that $g \circ \pi |\psi\rangle = |\psi\rangle$.*

This is sufficient to remove single-qubit gates in situations where we would otherwise use an eigenstate.

For many single-qubit gates, there is a trivial stabilizer state which stabilizes it. For instance, X is stabilized by $|+\rangle$, R_Z is stabilized by $|0\rangle$, and many other single-qubit stabilizer gates are conjugate to one of these cases. Now consider the gate θ_{X+Z}, whose eigenstates (unnormalized) $(1 \pm \sqrt{2})|0\rangle + |1\rangle$ are *not* stabilizer states. How then, given the gate $\theta_{X+Z} \otimes \theta_{X+Z}$, does one generate the gate θ_{X+Z} which acts only on one qubit? Han-Hsuan Lin discovered the first explicit nine qubit stabilizer state for this task.

Let π be a circuit that cyclicly permutes qubits 2 through 9, and suppose θ_{X+Z} is applied to qubit 1. Let $|\psi\rangle$ be the state stabilized by the following commuting Pauli strings,

$$\begin{array}{llll} XXZXZIIII, & ZIXZXZIII, & XIIXZXZII, & ZIIIXZXZI, \\ XIIIIXZXZ, & ZZIIIIXZX, & XXZIIIIXZ, & ZZXZIIIIX, \\ YYIIIYIII, & -YIYIIIYII, & YIIYIIIYI, & -YIIIYIIIY, \end{array}$$

9 of which are independent. One can check that conjugating each generator by $\theta_{X+Z} \circ \pi$ yields another element of the stabilizer group, so $(\theta_{X+Z} \circ \pi)|\psi\rangle = |\psi\rangle$. In other words, Conjecture 21 holds for θ_{X+Z}, and for all conjugates θ_{P+Q} by symmetry.

All that remains to verify the conjecture is to find a similar state stabilizing the eight remaining gates—the Γ gate and its conjugates. It suffices to find a stabilizer state $|\psi\rangle$ and circuit C, constructed of SWAP gates and a single Γ gate, such that $C|\psi\rangle = |\psi\rangle$.

12 Acknowledgments

We would like to thank Scott Aaronson for his insightful guidance throughout this project. Also, thanks to Han-Hsuan Lin and Adam Bouland for useful discussions.

References

[1] Scott Aaronson and Adam Bouland. Generation of universal linear optics by any beam splitter. *Phys. Rev. A*, 89(6):062316, 2014. arXiv:1310.6718.

[2] Scott Aaronson and Daniel Gottesman. Improved simulation of stabilizer circuits. *Phys. Rev. A*, 70(052328), 2004. arXiv:quant-ph/0406196.

[3] Scott Aaronson, Daniel Grier, and Luke Schaeffer. The classification of reversible bit operations. *ITCS*, 2017. arXiv:1504.05155.

[4] Panos Aliferis. Level reduction and the quantum threshold theorem. *arXiv preprint quant-ph/0703230*, 2007.

[5] Jonas T. Anderson. On the power of reusable magic states. *arXiv preprint arXiv:1205.0289*, 2012.

[6] A. Barenco, C. H. Bennett, R. Cleve, D. P. DiVincenzo, N. Margolus, P. Shor, T. Sleator, J. Smolin, and H. Weinfurter. Elementary gates for quantum computation. *Phys. Rev. A*, 52(3457), 1995. arXiv:quant-ph/9503016.

[7] Adam Bouland, Laura Mancinska, and Xue Zhang. Complexity classification of two-qubit commuting Hamiltonians. *arXiv preprint arXiv:1602.04145*, 2015.

[8] Sergey Bravyi and Alexei Kitaev. Universal quantum computation with ideal clifford gates and noisy ancillas. *Physical Review A*, 71(2):022316, 2005.

[9] Andrew M. Childs, Debbie Leung, Laura Mančinska, and Maris Ozols. Characterization of universal two-qubit Hamiltonians. *arXiv preprint arXiv:1004.1645*, 2010.

[10] John Horton Conway, Robert T. Curtis, Simon P. Norton, and Richard A. Parker. *Atlas of finite groups*. Oxford University Press, 1985.

[11] Jeroen Dehaene and Bart De Moor. Clifford group, stabilizer states, and linear and quadratic operations over GF(2). *Physical Review A*, 68(4):042318, 2003.

[12] Daniel Gottesman. Stabilizer codes and quantum error correction. *arXiv preprint quant-ph/9705052*, 1997.

[13] Daniel Gottesman. The Heisenberg representation of quantum computers. *arXiv preprint quant-ph/9807006*, 1998.

[14] Daniel Gottesman. Theory of fault-tolerant quantum computation. *Physical Review A*, 57(1):127, 1998.

[15] Walter Grimus and Patrick Otto Ludl. On the characterization of the SU(3)-subgroups of type C and D. *Journal of Physics A: Mathematical and Theoretical*, 47(7):075202, 2014.

[16] N. Cody Jones, Rodney Van Meter, Austin G. Fowler, Peter L. McMahon, Jungsang Kim, Thaddeus D. Ladd, and Yoshihisa Yamamoto. Layered architecture for quantum computing. *Physical Review X*, 2(3):031007, 2012.

[17] Martin Van den Nest. Classical simulation of quantum computation, the Gottesman-Knill theorem, and slightly beyond. *arXiv preprint arXiv:0811.0898*, 2008.

[18] R. Raussendorf, D. E. Browne, and H. J. Briegel. Measurement-based quantum computation on cluster states. *Phys. Rev. A*, 68, 2003. quant-ph/0301052.

[19] Yaoyun Shi. Both Toffoli and controlled-NOT need little help to do universal quantum computation. *Quantum Information and Computation*, 3(1):84–92, 2002. quant-ph/0205115.

[20] Peter W. Shor. Fault-tolerant quantum computation. In *Foundations of Computer Science, 1996. Proceedings., 37th Annual Symposium on*, pages 56–65. IEEE, 1996.

[21] Andrew Steane. Multiple particle interference and quantum error correction. *Proc. Roy. Soc. London*, A452:2551–2577, 1996. quant-ph/9601029.

A Enumeration

Theorem 22. *Let $\#\langle\cdot\rangle_n$ denote the number of n qubit gates in a class. Then*

$$\begin{aligned}
\#\langle G\rangle_n &= |G|^n\, n! && \text{for } G \text{ a group of single-qubit gates,}\\
\#\langle C(Z,Z), G\rangle_n &= |G|^n\, 2^{n(n-1)/2} n! && \text{for } \langle Z\rangle \subseteq G \subseteq \langle \mathcal{P}, \mathrm{R}_Z\rangle \text{ a group,}\\
\#\langle C(Z,X), \mathcal{P}\rangle_n &= 4^n 2^{n(n-1)/2} \prod_{i=1}^{n} (2^i - 1),\\
\#\langle C(Z,X), \mathcal{P}, \mathrm{R}_Z\rangle_n &= 8^n 2^{n(n-1)} \prod_{i=1}^{n} (2^i - 1),\\
\#\langle \mathrm{T}_4, \mathcal{P}\rangle_n &= 4^n a(n),\\
\#\langle \mathrm{T}_4, \mathcal{P}, \mathrm{R}_Z\rangle_n &= 8^n 2^{n(n-1)/2} a(n),\\
\#\langle \mathrm{T}_4, \mathcal{P}, \Gamma\rangle_n &= 4^n 2^{n(n-1)/2} \prod_{i=1}^{n} (2^i - (-1)^i),\\
\#\langle \mathsf{ALL}\rangle_n &= 4^n 2^{n^2} \prod_{i=1}^{n} (4^i - 1),
\end{aligned}$$

where

$$a(n) = \begin{cases} 2^{m^2} \prod_{i=1}^{m-1} (2^{2i} - 1), & \text{if } n = 2m,\\ 2^{m^2} \prod_{i=1}^{m} (2^{2i} - 1), & \text{if } n = 2m+1. \end{cases}$$

Proof. Most of these numbers follow from the lemmas above. For example, consider the class $\langle C(Z,X), \mathcal{P}, \mathrm{R}_Z\rangle$. It follows from Lemma 8 that any gate in this class has a circuit consisting of a layer of $C(Z,X)$ gates, then a layer of $C(Z,Z)$ gates, then a layer of single-qubit gates in G.

We would like to count the number of possible gates by multiplying the number of possibilities for each layer, but we must be careful that there is no gate with two circuit representations. Suppose for a contradiction that g_1 and g_2 generate the same gate, but some layer of g_1 differs from g_2. Then $g_1^{-1} g_2$ is the identity, since g_1 and g_2 generate the same transformation.

On the other hand, the $C(Z,X)$ layers of g_1 and g_2 meet in the middle of the circuit for $g_1^{-1} g_2$. If those layers do not generate the same linear transformation, then the combination is some non-trivial linear transformation which is, in particular, not Z-degenerate. The other layers of g_1 and g_2 *are* Z-degenerate, so we conclude that $g_1^{-1} g_2$ is not Z-degenerate (if it were, we could invert the outer layers to show that the two middle layers are Z-degenerate). But $g_1^{-1} g_2 = I$ is clearly Z-degenerate, therefore the $C(Z,X)$ layers of g_1 and g_2 must generate the same linear transformation.

The $C(Z,X)$ layers of g_1 and g_2 cancel (since we have shown they are equivalent), so they effectively disappear, and we make a similar argument about the $C(Z,Z)$ layers, and then the single-qubit layers. That is, if the $C(Z,Z)$ layers do not contain the same set of $C(Z,Z)$ gates, then we obtain a contradiction because they produce a non-degenerate layer in the middle, implying that $g_1^{-1} g_2 = I$ is non-degenerate. Once we remove the $C(Z,Z)$ layers, the single-qubit layers must be the same or they would leave behind a non-trivial single-qubit gate. We conclude that all layers of g_1 and g_2 are actually the same, so the number of gates is the product of the number of choices for each layer.

Now the problem is to count the number of choices for each layer. For the single qubit layer, this is clearly just n independent choices of single-qubit gate from $\langle \mathcal{P}, \mathrm{R}_Z \rangle$, or 8^n. For the $C(Z,Z)$ layer, there is a choice whether or not to place a $C(Z,Z)$ gate in each of the $\binom{n}{2}$ possible positions, so $2^{n(n-1)/2}$ choices for the layer. For the $C(Z,X)$ layer, observe that $C(Z,X)$ generate precisely the set of invertible linear transformations, of which there are

$$2^{n(n-1)/2} \prod_{i=1}^{n} (2^i - 1)$$

by a classical argument. Multiplying the three layers, we have a total of

$$\# \langle C(Z,X), \mathcal{P}, \mathrm{R}_Z \rangle_n = 8^n 2^{n(n-1)} \prod_{i=1}^{n} (2^i - 1)$$

n-qubit transformations generated by $C(Z,X)$, $\mathcal{P}$, and R_Z.

The numbers for $\langle G \rangle$, $\langle C(Z,Z), G \rangle$, $\langle C(Z,X), \mathcal{P} \rangle$, $\langle \mathrm{T}_4, \mathcal{P} \rangle$, and $\langle \mathrm{T}_4, \mathcal{P}, \mathrm{R}_Z \rangle$ follow by a similar argument, although for the last two classes we need the fact that T_4 generates

$$a(n) = \begin{cases} 2^{m^2} \prod_{i=1}^{m-1} (2^{2i} - 1), & \text{if } n = 2m, \\ 2^{m^2} \prod_{i=1}^{m} (2^{2i} - 1), & \text{if } n = 2m+1. \end{cases}$$

orthogonal transformations on n qubits.

For the final two classes, we use known expressions (from [10]) for the number of $n \times n$ unitary matrices over $\mathbb{F}_4$ (in the case of $\langle \mathrm{T}_4, \mathcal{P}, \Gamma \rangle$) and for the number of $2n \times 2n$ symplectic matrices over $\mathbb{F}_2$ (in the case of ALL). We multiply by 4^n in both cases to account for the phase bits, which are completely independent of the matrix part. □

Theorem 23. *The asymptotic size of each class is as follows.*

$$\begin{aligned}
\log_2 \# \langle G \rangle_n &= n \log_2(|G|) + n \log_2 n - n \log_2 e + \frac{1}{2} \log_2 2\pi + O\left(\frac{1}{n}\right), \\
\log_2 \# \langle C(Z,Z), G \rangle_n &= n \log_2(|G|) + \frac{n(n-1)}{2} + n \log_2 n - n \log_2 e + \frac{1}{2} \log_2 2\pi + O\left(\frac{1}{n}\right), \\
\log_2 \# \langle C(Z,X), \mathcal{P} \rangle_n &= n^2 + 2n - \alpha + O(2^{-n}), \\
\log_2 \# \langle C(Z,X), \mathcal{P}, \mathrm{R}_Z \rangle_n &= \frac{3}{2} n^2 + \frac{5}{2} n - \alpha + O(2^{-n}), \\
\log_2 \# \langle \mathrm{T}_4, \mathcal{P} \rangle_n &= \frac{1}{2} n^2 + \frac{3}{2} n - \beta + O(2^{-n}), \\
\log_2 \# \langle \mathrm{T}_4, \mathcal{P}, \mathrm{R}_Z \rangle_n &= n^2 + 3n - \beta + O(2^{-n}), \\
\log_2 \# \langle \mathrm{T}_4, \mathcal{P}, \Gamma \rangle_n &= n^2 + 2n + \gamma + O(2^{-n}), \\
\log_2 \# \langle \mathsf{ALL} \rangle_n &= 2n^2 + 3n - \beta + O(4^{-n}).
\end{aligned}$$

where G is the same as in Theorem 22, and

$$\alpha = -\sum_{i=1}^{\infty} \log_2(1 - 2^{-i}) \approx 1.7919,$$
$$\beta = -\sum_{i=1}^{\infty} \log_2(1 - 4^{-i}) \approx 0.53839,$$
$$\gamma = \sum_{i=1}^{\infty} \log_2(1 - (-2)^{-i}) \approx 0.27587.$$

Proof. We take the logarithm of each class size, which we can separate into the logarithm of each layer comprising that class, as in Theorem 22. For most layers this is straightforward, except for the layer of permutations, orthogonal transformations, or general linear transformations. The first we handle with Stirling's approximation. For the other two, we factor out powers of two leaving a partial sum of a convergent series, which we analyze with a Taylor expansion. The classes $\langle \mathrm{T}_4, \mathcal{P}, \Gamma \rangle_n$ and $\langle \mathsf{ALL} \rangle_n$ follow by similar techniques. □

Corollary 24. *Let $\mathcal{C}$ be any class, and let G be an n-qubit gate chosen uniformly at random from $\mathcal{C}$. Then*

$$\Pr[G \text{ generates } \mathcal{C}] = 1 - O(2^{-n}).$$

B Classical reversible gates with quantum ancillas

In this section we describe what the classical reversible gate lattice of Aaronson et al. [3] would have looked like had they allowed quantum rather than classical ancillas. We extend all classical gates discussed in that paper to the quantum setting in the most natural way; that is, a classical reversible gate acts as a permutation on the computational basis states. Figure 11 shows this dramatically simpler lattice.

Some of the collapses in the lattice are immediate. For instance, the class $\langle \mathrm{NOT} \otimes \mathrm{NOT} \rangle$ collapses with the class $\langle \mathrm{NOT} \rangle$ because $\mathrm{NOT}\,|+\rangle = |+\rangle$. A similar collapse occurs been all classes where the parity issue arises such as between the classes $\langle \mathrm{CNOTNOT} \rangle$ and $\langle \mathrm{CNOT} \rangle$.

A more interesting collapse occurs between all mod-k-preserving classes for $k \geq 2$. Consider the following gate $G : \{0,1\}^k \to \{0,1\}^k$ of order 2 which preserves Hamming weight mod k:

$$G(0^k) = 1^k$$
$$G(1^k) = 0^k$$
$$G(1^a 0^b 0) = 1^{a-1} 0^{b+1} 1$$
$$G(1^a 0^b 1) = 1^{a+1} 0^{b-1} 0$$

where G acts as the identity on all other inputs. Since G preserves the Hamming weight mod k, it must appear in the class. We will show that G can generate a NOT gate. Let

$$|\psi_k\rangle := \frac{1}{\sqrt{k}} \sum_{i=0}^{k-1} \left| 1^i 0^{k-i-1} \right\rangle$$

so, for example

$$|\psi_4\rangle = \frac{|000\rangle + |100\rangle + |110\rangle + |111\rangle}{2}.$$

Now, for $b \in \{0,1\}$, $G(|\psi_k\rangle|b\rangle) = |\psi_k\rangle|b \oplus 1\rangle$. Therefore, each mod-$k$-preserving class for $k \geq 2$ collapses to the $\langle \text{Fredkin}, \text{NOT} \rangle$ class. Furthermore, Figure 10 shows that the Fredkin and NOT gates are sufficient to generate a CNOT. Therefore, every non-conservative non-affine class generates all classical reversible transformations.

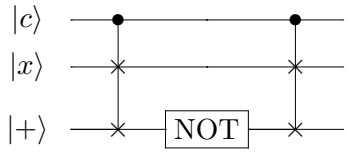

Figure 10: Generating CNOT from Fredkin and NOT gates

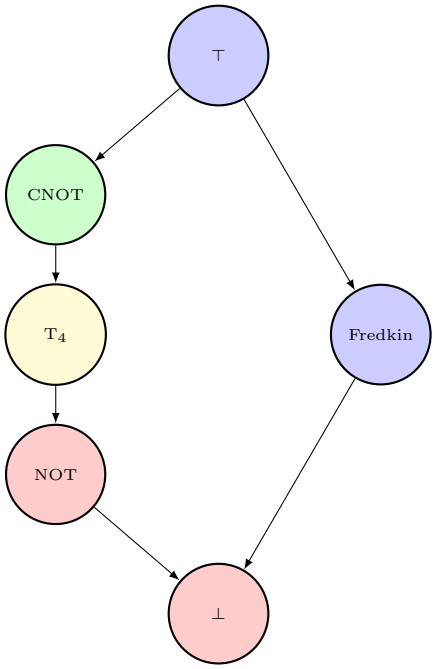

Figure 11: The inclusion lattice of classical gates using quantum ancillas

We now only need to prove that the classes appearing in Figure 11 are distinct. Notice, however, that the classes $\langle \text{CNOT} \rangle$, $\langle \text{T}_4 \rangle$, and $\langle \text{NOT} \rangle$ all have stabilizer generators, which by the results of this paper, generate distinct classes. We only need to show then, that the $\langle \text{Fredkin} \rangle$ class is distinct from the remaining classes. However, the invariant in [3] more or less functions to prove this separation. Namely, Fredkin conserves the Hamming weight of its input. Therefore the sum of the Hamming weights of the computational basis states of the input state is conserved. However, the NOT gate necessarily changes this sum, witnessing that $\text{NOT} \notin \langle \text{Fredkin} \rangle$, and therefore that the lattice is complete.

C Three-qubit generator for $\langle \mathrm{T}_4, \Gamma, \mathcal{P} \rangle$

In this section, we point out a peculiar consequence of the classification. Namely, that there exists a class whose smallest generator is on 3 qubits. Contrasted with the affine gate sets of Aaronson et al. [3] which have minimal generators over 1, 2, 4, and 6 bits. The existence of such a class is an immediate consequence of the enumeration in Appendix A. That is,

$$\begin{aligned}
&\#\langle \mathrm{T}_4, \mathcal{P}, \Gamma \rangle_3 = 2^{3(3-1)/2+2(3)} \prod_{i=1}^{3} (2^i - (-1)^i) = 41472 \\
&\#\langle \Gamma, \mathcal{P} \rangle_3 = 12^3 3! = 10368 \\
&\#\langle \mathrm{T}_4, \mathcal{P} \rangle_3 = 4^3 2^{1^2} \prod_{i=1}^{1} (2^{2i} - 1) = 384 \\
&\#\langle \mathcal{P} \rangle_3 = 4^3 3! = 384
\end{aligned}$$

so the non-degenerate gates in $\langle \mathrm{T}_4, \mathcal{P}, \Gamma \rangle_3$ outnumber the degenerate gates (i.e. $\#\langle \Gamma, \mathcal{P} \rangle_3$), contrasted with the class $\langle \mathrm{T}_4, \mathcal{P} \rangle$, where this does not occur. Notice that there are 4 cosets of $\langle \Gamma, \mathcal{P} \rangle_3$ in $\langle \mathrm{T}_4, \mathcal{P}, \Gamma \rangle_3$ by Lagrange's Theorem, corresponding to 4 gates that are nonequivalent up to applications of elements in $\langle \Gamma, \mathcal{P} \rangle_3$. If we let $\alpha = \left(\begin{smallmatrix} 0 & 1 \\ 1 & 1 \end{smallmatrix}\right) \in R$, then one such gate is described by the following tableau

$$\left(\begin{array}{ccc|c} \alpha & I & I & 0 \\ I & \alpha & I & 0 \\ I & I & \alpha & 0 \end{array}\right).$$

By Theorem 19 it is clear that this gate indeed generates all of $\langle \mathrm{T}_4, \mathcal{P}, \Gamma \rangle$.

D Canonical form for 2-qubit circuits

In this section, we describe a very clean canonical form for 2-qubit stabilizer circuits.

Theorem 25. *Let C be any stabilizer circuit on two qubits. Then, C is equivalent to a circuit of at most depth 3 composed of the following sequence of gates*

1. *a SWAP gate, and*
2. *a tensor product of single-qubit gates, and*
3. *a generalized CNOT gate,*

where we can choose at each step whether or not to include the gate. That is, C is of the form of the circuit depicted in Figure 12.

Proof. Since C is a stabilizer circuit, it can be written as a product of CNOT, θ_{X+Z}, and R_Z gates.[7] Recall that conjugating a generalized CNOT gate by a single-qubit gate is simply another generalized CNOT gate. Therefore, we can push all the single-qubit gates left and all the generalized CNOT gates right. All that remains to show is that we can coalesce the generalized CNOT gates into a single CNOT gate. We refer to Table 5 for those equivalences, and note that identical

[7]Notice that we don't need to include SWAP gates because of the equivalence $\mathrm{SWAP}(1,2) = \mathrm{CNOT}(1,2) \circ \mathrm{CNOT}(2,1) \circ \mathrm{CNOT}(1,2)$.

generalized CNOT gates cancel. Eventually, what remains is a circuit composed of single-qubit gates, SWAP gates, and at most one generalized CNOT gate. We can push the SWAP gates to the left (they collapse to either a single SWAP gate or the identity) and combine the single-qubit gates, which completes the proof. □

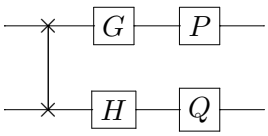

Figure 12: Canonical form of a 2-qubit circuit: optional SWAP gate, optional $C(P, Q)$ gate, and single-qubit gates G and H

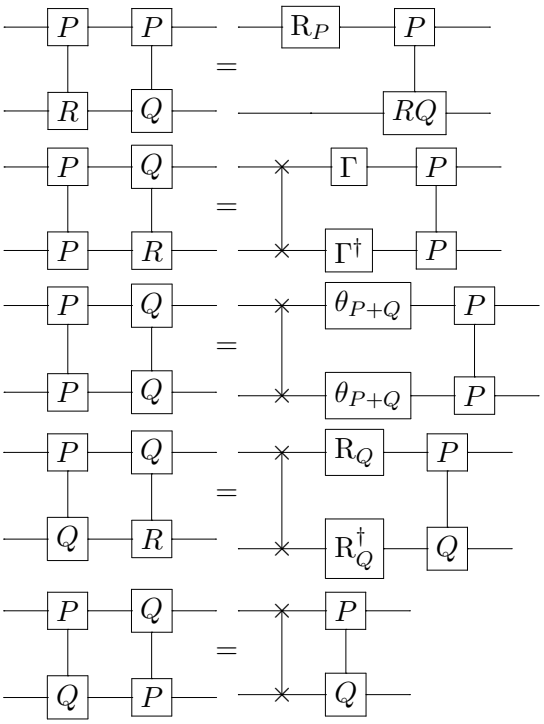

Table 5: Rules for coalescing generalized CNOT gates, assuming $\Gamma P \Gamma^\dagger = Q$ and $\Gamma Q \Gamma^\dagger = R$